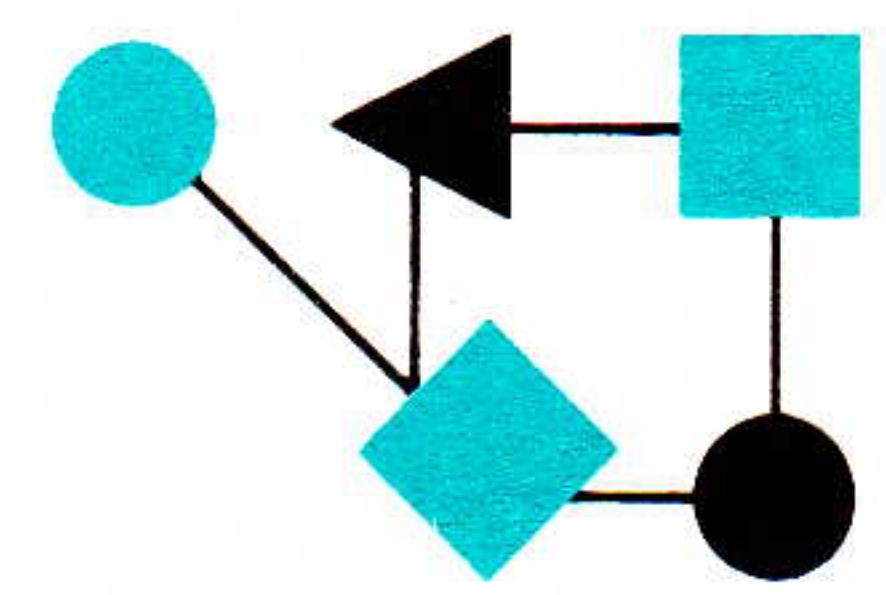


CONNEXIONS



The Interoperability Report

November 1990

Volume 4, No. 11

*Connexions—
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

FTAM, FTP, and NFS.....	2
Network Management Directions.....	10
INTEROP 90 wrap-up.....	13
Profile: NORDUnet.....	18
SIGCOMM Award winners...	23
Press here for the Internet...	24
Book Reviews.....	26
Upcoming Events.....	29
GOSIP documents.....	30

Connexions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779.

Copyright © 1990 by Interop, Inc.
Quotation with attribution encouraged.

Connexions—The Interoperability Report
and the *Connexions* masthead are
trademarks of Interop, Inc.

ISSN 0894-5926

From the Editor

For me, INTEROP 90 started in the afternoon of Saturday, October 5 when I attended the briefing for the "network team," the 75 plus people who had volunteered for the job of building a large internet in the space of only a few hours. Later that night the project got underway and by noon Sunday, over 25 miles of cable was in place ready to be connected to the more than 200 vendor booths.

But INTEROP is more than just an exhibition of state-of-the-art networking technology and many very impressive interoperability demonstrations. There were also 22 tutorials, 45 conference sessions, vendor presentation theaters, and dozens of informal Birds Of a Feather (BOF) sessions.

This issue will look at INTEROP 90 from the perspective of an attendee, our assigned correspondent Daniel Dern who tried to see and do as much as possible during this very hectic week. His report is followed by a few snapshots from this year's show. We also feature several articles that relate to INTEROP in one way or another. First out is a comparison of FTAM, FTP and NFS for Enterprise Networks. The article is by Eric Fleischman who chaired a very well attended BOF on this topic at INTEROP 90.

Network Management was discussed in several tutorials, conference sessions and BOFs at INTEROP. Also featured on the exhibit floor was the largest cooperative demonstration to date of the *Simple Network Management Protocol* (SNMP). Karl Auerbach and Denis Yaro reflect on the future of Network Management in an article starting on page 10.

INTEROP is becoming an international conference, in terms of the attendees, speakers and topics covered. This year's conference featured no less than ten international speakers, most notably several representatives from the Soviet Union and Eastern Europe. The changing political climate which allowed them to attend is also opening the way for increased international cooperation in the area of computer networking. *Connexions* has featured profiles of large internets in the past. This time we take an international perspective and look at an area of the world dear to the Editor's heart, namely Scandinavia. Mats Brunell describes the NORDUnet network in an article on page 18.

The TCP/IP Internet is now so large (and still rapidly growing) that it deserves more than just a casual mention in the trade press. How to best achieve such coverage was the topic for another BOF at INTEROP 90. We asked Daniel Dern to summarize the discussions.

OSI/GOSIP was also a theme for INTEROP 90. We will cover this topic in future issues. This month we bring you information on how to order the relevant GOSIP documents.

FTAM, FTP, and NFS Within an Enterprise Network

by Eric Fleischman, Boeing Computer Services

Enterprise Networks

Many large corporations are automating their internal processes by combining their disparate network pieces into a single, unified network which unites their company and connects to their major suppliers and customers. This resulting "global" network, called an *Enterprise Network*, may contain tens (or hundreds) of thousands of nodes. The establishment of an Enterprise Network presents a daunting challenge to the network designer. The challenges arise from the necessity of having all of the different network pieces be able to "plug and play" at each interconnection juncture if business activities are to be successfully automated. Interconnection junctures occur at several levels of communication: the physical media, the subnetwork, the network transport, the network application and the data interchange format.

Contenders

This article is concerned with the application layer technology which permits distributed file sharing across a network. Distributed file sharing technology is implemented by file transfer or file access protocols. An axiomatic assumption of this paper is that a common protocol for distributed file sharing is needed within a network in order to maximize interconnectivity and flexibility. Currently there are three primary contenders for this role of key distributed file sharing protocol within an Enterprise Network. The contenders are:

- *File Transfer, Access and Management* (FTAM)
- *File Transfer Protocol* (FTP)
- *Network File System* (NFS).

The goal of this article is to evaluate the suitability of each of these protocols within a large Enterprise Network. Due to space requirements, this article is a very condensed version of a substantially larger and more authoritative paper, [1], which readers should consult for a fuller treatment of this topic.

There are three possible components to distributed file sharing protocols: *file access*, *file transfer*, and *file management*. File access refers to the reading, writing or deleting of selected parts (e.g., records) of a file residing on a remote end-system (i.e., a network node). File transfer, on the other hand, refers to the movement of complete files between two end-systems. File management refers to remotely reading or altering file attributes (e.g., filename) that define a file within the operating system within which it resides. It is usually the case that file management capabilities are viewed as appendages to file access or file transfer protocols.

FTAM

FTAM is an OSI international standard defined by the International Organization for Standardization (ISO) in the ISO 8571 documents [2]. It has strong file access, file transfer, and file management capabilities. Implementations of FTAM have begun to appear for many different platforms. It is anticipated that the "[OSI] protocols will play a major role in worldwide communications because almost all countries in the world have made a major political commitment to the adoption of OSI" [3]. In the United States, FTAM became a Federal procurement requirement beginning August 15, 1990 as part of the US Government's GOSIP requirements [4]. In addition, it is currently a MAP/TOP industry requirement.

FTP	FTP was defined by the United States Department of Defense (DoD) in the MIL-STD-1780 document [5]. FTP is an integral part of the DoD standards which define the Internet, ARPANET, and other TCP/IP networks. As such, it is a popular protocol with a mature implementation base. FTP is primarily a file transfer protocol with weak file management capabilities.
NFS	NFS is a popular de facto standard originally developed by Sun Microsystems, Inc. It is primarily a file access protocol with UNIX-like file management capabilities. It has recently been submitted to the Internet as an RFC (<i>Request for Comments</i>) [6]. NFS is one of six public domain network services developed by Sun Microsystems. These services are part of Sun's <i>Open Network Computing</i> (ONC) model. "Today more than 290 organizations, including 150 computer manufacturers and leading software firms, have licensed ONC/NFS to use in their operating systems" [7].
Comparison of target domains	Each of the three protocols has different target domains, different design goals and philosophies, and different strengths and weaknesses. The following chart displays the probable target domain for each approach: <p>FTAM: Permit users to share file data regardless of the underlying hardware or software differences between computers in an OSI environment.</p> <p>FTP: Provide a simple solution for transferring whole files in a TCP/IP environment.</p> <p>NFS: Provide a seamless extension of the UNIX filesystem so that UNIX (or UNIX-like) computers may be able to transparently share data in a workgroup environment.</p> <p>NFS is file access oriented. It streamlines one's ability to access parts of files, while at the same time supporting file transfer and file management. FTP, on the other hand, is primarily a file transfer protocol with essentially no file access abilities and very limited file management capabilities. FTAM sits on the fence between the two in that it provides substantial support for file access, file transfer, and file management. Current implementations of FTAM, however, have yet to implement the full richness of the FTAM protocol's functionality. For example, current FTAM implementations are significantly clumsier with less functionality than NFS in the file access domain.</p>
Key design differences	The design orientation of NFS differs considerably from both FTAM and FTP. FTAM and FTP were originally designed to be data communications protocols without regards to the underlying operating system on the end systems. NFS, by contrast, is a distributed operating system extension. Much like UNIX System V's <i>Remote File System</i> (RFS) and Apollo's <i>Network Computer System</i> (NCS), the original goal of NFS was to extend a UNIX filesystem across a distributed workgroup environment. From this beginning, NFS has evolved into a system which functions in a diversity of hardware and software platforms within a much larger domain than merely a local LAN segment. However, the original orientation has enduring influence within the protocol as a whole.
Stateless protocol	NFS is a stateless protocol in which client requests (the "user") are bundled into Sun Microsystems' <i>Remote Procedure Call</i> (RPC) packets and sent as datagrams to the servers which will process the requests. "Stateless," in this context, means that the server has no memory of previous RPC requests—each request stands on its own merit.

continued on next page

FTAM, FTP, and NFS (*continued*)

The protocol is synchronous in the sense that a user process will block until the server response is obtained. Because a connectionless, confirmationless, datagram service (UDP/IP) is used to transmit the packets, the client assumes that the packet has been lost if no response from the server has been obtained within a configurable time period. Once the time period expires ("timeouts") without a response from the server, the client retransmits the packet. This occurs a configurable number of iterations. Should a non-responding server's files have been originally "soft mounted," the remote procedure call will eventually fail and the user will be alerted to the failure. Should the remote files have been "hard mounted," however, then the retransmission cycle will continue until such a time as the server finally responds.

FTAM and FTP, on the other hand, rely upon a connection oriented transport service to provide a confirmed service. This is needed because these protocols contain distinct internal states and the confirmed service permits both end systems to track both their own state as well as the state of their peer.

File translation abilities

Enterprise Networks frequently contain a diverse, heterogeneous mix of operating systems and files within their end systems. The concept of "file" means different things to different operating systems. For example, for many operating systems (e.g., UNIX, MS-DOS, and CP/M) a file consists of a sequential ordering of characters. These characters are traditionally viewed as comprising either text (e.g., ASCII) or binary (e.g., executable or object code, bit maps, etc.) files. By contrast, other operating systems (such as DEC's VMS and IBM's MVS and VM) view files as being composed of a series of records or pages. In this approach, the records on a file may be organized in a variety of ways. The most popular file organization schemes for this approach are: *sequential* (records are placed in physical order); *indexed sequential* (records are arranged in logical sequence according to a key contained in each record); *direct* (records are randomly accessed by their physical addresses on a *Direct Access Storage Device* [DASD]); and *partitioned* (a file is composed of sequential sub-files). Thus, it is important to realize that a file is more than merely a collection of data—it also contains information which permits a specific operating system to organize and understand that data. This "organizational component" of a file is operating system specific: different operating systems will most likely not be able to understand the organizational components of other operating systems and thus will not be able to "understand" foreign files' contents.

FTAM is designed to recognize that a great many types of files exist. It is general enough to have provisions to define (via common agreement between both end-systems) new file types so that the file structure will be correctly translated. The OSI protocols themselves hide file conversion concerns from the user. That is, if a file is copied via FTAM between diverse computers, the file data is transferred but the operating system specific aspects are translated into the equivalent on the destination computer. For example, a file copied from an MVS computer to a UNIX computer by FTAM will be converted into the proper UNIX format in order to preserve the original MVS data and semantics.

FTP has provision to support the conversion of the most common types of files between end systems. This definition should be adequate for a large percentage of the file types in common usage.

However, the more esoteric types of files will be excluded. Thus, FTP has a subset of the capability of FTAM in this regards.

NFS, on the other hand, has no built-in file conversion capabilities. NFS copies the “raw” file bit stream without translating any of the file data. The net result is that the file appears to the client with the same bit-stream which it had on the server—regardless of whether the client and server would interpret that bit stream in the same way or not. Since the semantics of the bit streams vary depending on the local operating system, a very real possibility exists for the destination end system to misunderstand the received data—unless, of course, the users have a prior knowledge of the received file’s contents and operating system differences so that they might use a utility to convert the file into the local equivalent. This lack of definition introduces opportunities for potentially grave interoperability problems.

[An aside: Informed readers will recognize that NFS’ RPC uses a data description and encoding standard named the *External Data Representation* (XDR) to transfer its header information between end systems. This encoding mechanism helps to insure that the remote node will understand the enclosed packet header information. However, XDR is not used to translate the data as is commonly assumed.]

Lastly, both NFS and FTP have an “English-language” bias. That is, NFS’ XDR only specifies ASCII characters and FTP solely supports ASCII and EBCDIC (as well as binary transfer). Both ASCII and EBCDIC encodings were designed for English and work quite well with it. However, they are inadequate for many other languages. FTAM avoids this problem by letting the specific encoding be part of the data presentation context negotiation. In addition, the international nature of OSI has purged such “ethnocentrism” from its standards.

Comparison of data translation

Data translation refers to preserving the semantics of the transferred data between end systems with different hardware orientations. This is tremendously important in a large, heterogeneous networking environment because computer data is context sensitive: a sequence of bits has different significance on different computers. The reason for this is fairly obvious: different types of computers have potentially different sized byte and word sizes—in fact, potentially different sizes for all entities—so that an untranslated bit stream will potentially be incorrectly partitioned.

For example, most computers have octet-sized bytes, but computers with six bit and nine bit bytes also exist. Similarly, word size varies [8]: The most common word sizes are 16, 32, or 64 bits long. However, the now obsoleted CDC 160 and the IBM 650 have 12 bit word sizes; large Univac and Honeywell computers commonly use 36-bit words; Burroughs’ larger systems use 48-bit words; and Control Data’s Cyber systems are 60-bit word machines. “Computers have [even] been built in which there was no preferred or fixed word length. Most of these were data processing machines that operated on character strings” [9]. Examples of the variable word length machines are the IBM 1400 series and the IBM 1620.

Additional complications arise in that even if two computers have the same sized entities, they may differ by being big- or little-endian (i.e., whether the high order bit is first or last); having different unit size alignment (i.e., 1, 2, 4, and 8 octet alignment units are all common); and whether their character strings are big- or little-endian (i.e., whether the high byte is first or last).

FTAM, FTP, and NFS (*continued*)

Consequently, the significance of data translation is that received untranslated data bit streams will only be intelligible to computers with a similar hardware base. If this is not the case, then the user must either convert the received data into the local equivalent or else face the real possibility that the received data is being misinterpreted. However, if the file transfer or file access protocol performs the data translation, the user is saved from this possibility.

Each of the three protocols have unique approaches as to how data translation is to be performed:

FTAM: Both end systems negotiate the data presentation context. Since the receiving computer consequently knows the semantics of the received bitstream, it will correctly translate the data into the local equivalent. All hardware and software platforms are thus supported.

FTP: Provide translations between 32 bit words and 36 bit words. The TELNET control channel will direct the form by which this transformation will occur over the data channel.

NFS: No data translation is done by the protocol. All data translation must be done by the users themselves.

One can see that the capabilities of FTAM in the data translation domain is substantially greater than that of either NFS or FTP. This is all the more significant because there is no provision (other than prior knowledge) within NFS or FTP for the discovery of what the data format of the remote computer may be. Thus, to confidently use either of these protocols in heterogeneous computer environments, one must have prior knowledge of each of the end systems to know whether (additional) data translation will be needed or not. If data translation is needed, the users will need to perform this translation themselves [10, 11]. One can well imagine that this additional overhead will restrict the number and type of divergent computers (and files) with which any given FTP or NFS user will be able to interoperate in a large heterogeneous networking environment. This is probably the most severe limitation of both NFS and FTP.

While NFS is outstanding in the environment for which it was designed, it has significant flaws when it is assigned the role of being the key file access protocol in an Enterprise Network. This is because NFS is not scalable. Scalability may be defined as the ability of a network to grow indefinitely without an upper bound. NFS has at least seven problems which effect its performance within an Enterprise Network. (Similar problems do not occur with FTAM and FTP.) The goal of this section is to outline these flaws which cumulatively make NFS a poor choice to serve as a key protocol within an Enterprise Network.

The first problem is that an NFS server will “thrash” due to a large volume of requests. As was previously mentioned, both FTAM and FTP use *connection-oriented* transport services. This means that if their servers are too busy to accept a new request, they will simply refuse the new transport connection and the client will have to try again later. In NFS, however, there is no provision for a busy NFS server to refuse a new request. The bottleneck of any server is its disk I/O which limits the rate at which it can perform file requests. Once the number of requests increases at a rate beyond which the server’s disk(s) can perform, outstanding requests will start to accumulate.

NFS shortcomings in an Enterprise Network

Thrashing

If this accumulation rate greatly exceeds the disk throughput rate, requests will begin to timeout. Retransmitted requests will further burden the server. Once this occurs, the number of requests start to snowball even with a moderate number of new user requests. (Of course, a similar situation may also occur simply by having a large number of users simultaneously make requests of the same server.)

Eventually the buffer space within the server which stores these requests is exceeded. At this point, the server will do one of two things: it either will “lose” new requests or else it will repeatedly swap its expanding “request buffer area” between its disk and its RAM as it tries to keep up with the snowballing number of requests. In any case, once this occurs the server’s throughput is reduced until it becomes effectively paralyzed. It should be noted that such a situation is extremely unlikely to occur in the small workgroup environment for which NFS was designed. However, in an Enterprise Network where literally thousands of clients may potentially desire to access the same server, there is a high probability for this situation to occur.

RPC timeouts

A second problem effecting NFS’ scalability is that RPC’s timeout mechanism begins to get “frayed at the edges” as the network grows in size. As the number of hops increases between end systems, due to the presence of numerous intermediate systems (e.g., routers and gateways), the time for a packet to reach a remote server increases. A parallel situation occurs when NFS is utilized in a low speed—56 Kbps or less—WAN environment. In such environments, one faces the real possibility of incurring one or more timeouts before the answer from the server arrives. This means that there may be several requests and several answers for a single user request in transit at a given time. This scenario is full of potential for unsatisfactory performance and miscommunication.

For example, an NFS server, being stateless, has no mechanism to tell whether a given request has been performed yet or not: each retransmitted request is treated as a new request. Consequently, it is very conceivable that a server in such a networking environment may be inadvertently requested to make multiple writes in response to a single user request.

Directory mounting

A third problem is that NFS requires an a priori set-up stage built into its use. That is, files are only available to a user if their directory had previously been locally mounted. Mounting of directories (i.e., devices) is usually done by a systems administrator. This means that one either must know in advance exactly which remote devices one wishes to use, or else one needs to get the system administrator to mount the desired device once one decides to access those files. In any case, one must have a prior knowledge of which files one is interested in accessing so that they can be initially placed in an accessible state.

UNIX bias

A fourth problem is that not all NFS servers can also be NFS clients. NFS has made UNIX its de facto filesystem standard. It currently requires that all of its implementations treat files exactly as the UNIX operating system treats its files. Such an approach creates many difficulties for the porting of NFS into some non-UNIX environments (particularly those with “flat” filesystems) because the concept of what a “filesystem” is differs between operating systems. A consequence of this is that while some non-UNIX operating systems (e.g., IBM’s MVS) are able to write utilities which can emulate a UNIX-oriented server, they cannot similarly write utilities which can function as NFS clients.

FTAM, FTP, and NFS (*continued*)

Machines without NFS clients are in effect excluded from full participation in an NFS-oriented network. This establishes a distinction between the “haves” and the “have-nots.” It is certain that the “have-nots” will not be satisfied with a perpetual inability to initiate file sharing across the network.

Concurrent file access

A fifth problem is associated with *concurrent access* and *file locking*. This semaphore-like functionality is needed to permit files to remain in a consistent and known state when they are simultaneously accessed by many users. The problem arises because NFS is stateless but the concepts of concurrent access and file locking are inherently “state-full.” NFS clearly needs these capabilities because of its strong file access orientation. However, such a service is “contrary” to the statelessness of the protocol. (By contrast, FTAM is defined with native file locking abilities and FTP need not worry about this problem because it has no file access capabilities.) In order to get around this problem, Sun Microsystems has provided a separate RPC-based file locking utility to perform this functionality. However, a number of problems, including the following, have arisen with this utility:

“Because the NFS server maintains no locks between requests, and a write may span several RPC requests, two clients writing to the same remote file can receive intermixed data on long writes.” [12].

A similar problem also manifests itself when NFS is ported to operating systems which do not have stateless file systems. For example, IBM’s MVS requires that if an application opens a file for a write, that the file be allocated and locked so that it will not be accessible by other applications. Thus, in these environments NFS is constrained to be state-full when it is designed to be stateless. Such a situation is filled with opportunities for undesired results—not to mention the problems associated with the current MVS server’s “work around” which inserts a timer to arbitrarily determine how long a given user is using a given MVS file.

Multiple writes

A sixth problem is that the statelessness of NFS leaves open the possibility of undesired multiple writes. Consider, for example, the scenario in which a user is writing to a server and the server successfully executes the write, but right before the server is able to send an acknowledging message to the client, the server crashes. While such a scenario is admittedly rare, the protocol specifies that the client will “timeout” and retransmit the request. The server will not recognize that it had already performed the retransmitted request and it will perform the write once again. Hence, an undesired write will occur.

Flow control problems

Finally, a connectionless transport service does not have the built-in flow control advantages that connection-oriented transport services possess. That is, many critics have claimed that NFS is a “bad citizen” on the network because an NFS datagram packet is sent whenever a client desires a service regardless of the current state of the network. This subjects the network to sudden “spikes” of activity during which it may become overloaded. The greater the number of machines running NFS (or some other datagram-oriented service), the greater the potential severity of the spikes. Should the network actually become overloaded, it is not unreasonable for a router or gateway to drop a packet. If this should occur, the network will slow as packets are re-transmitted to recover the lost packets. This, in turn, will delay response times.

Conclusions

NFS has many scalability limitations which make it inadvisable to be selected as the key file sharing protocol within an Enterprise Network. In contrast, either FTAM or FTP would be good choices for that role. FTP's simplicity and existing user base makes it extremely attractive to be the key file sharing protocol within an Enterprise Network. This is particularly the case in the immediate short term when most FTAM implementations are immature and not yet widely deployed. However, because of FTP's lack of a file access capability and, more importantly, its incomplete data translation abilities, it is a substantially poorer long-term choice than FTAM.

Consequently, FTAM will increasingly be the best choice for a large, heterogeneous Enterprise Network. Its strengths lie in its international stature; its strong data translation abilities; and its balanced support for file access, file transfer, and file management. Since it is currently required by many governments (including the US Government) and a growing number of industry groups throughout the world, it will become an increasingly valuable tool in the world-wide marketplace. For these reasons, FTAM is the best choice to serve as a file transfer technology in a large Enterprise Network.

References

- [1] Eric Fleischman, "Network File Transfer and File Access Protocols: FTAM, FTP, NFS," Unpublished manuscript. Copies may be obtained from the author at: Boeing Computer Services, PO Box 24346, MS 7M-46, Seattle, WA 98124-0346.
- [2] ISO 8571, "Information Processing Systems—Open Systems Interconnection—File Transfer, Access and Management," parts 1–5.
- [3] Daniel C. Lynch, "TCP/IP and OSI," *Journal of Data & Computer Communications*, summer 1989.
- [4] "Government Open Systems Interconnection Profile (GOSIP)," FIPS PUB 146, August 24, 1988.
- [5] Military Standard: File Transfer Protocol, MIL-STD-1780, 1984.
- [6] "NFS: Network File System Protocol specification," Sun Microsystems, Inc., RFC 1094, 1989.
- [7] ONC/NFS Solution Guide. Second Edition. Sun Microsystems, Inc., September 1989, Pages 1–44.
- [8] Anthony Ralson and Edwin Reilly, Jr., "Encyclopedia of Computer Science and Engineering," Second Edition. Van Nostrand Reinhold Company, New York, ISBN 0-442-24496-7, 1983, Pages 13, and 1433–1434.
- [9] ibid. Anthony Ralson and Edwin Reilly, Jr., Page 1572.
- [10] ibid. MIL-STD-1780, page 9.
- [11] Personal conversation with Chuck McManis on Dec. 6, 1989. Chuck McManis works for Sun Microsystems, Inc. as a software engineer. He was part of the team which developed NFS.
- [12] *In Depth*, Sun Microsystems, Inc. July 1989, Page 931.

ERIC FLEISCHMAN is a Senior Data Communications Analyst in Boeing Computer Service's Architecture and Standards Group. He is part of the team which is designing the Boeing Enterprise Network. In addition, Eric is involved in activities external to Boeing including being the North American Representative to the MAP/TOP World Federation for OSI testing and a member of the COS Architectural Committee. Eric has previously worked at Victor Technologies, Digital Research Inc., and AT&T Bell Laboratories. He has degrees from Wheaton College (Illinois), University of Texas at Arlington, and University of California at Santa Cruz.

Network Management Directions

by
Karl Auerbach, Epilogue Technology
and
Denis Yaro, Sun Microsystems

Network Management is one of the most widely discussed and debated topics in the networking community. Although the subject has only recently become prominent, network management has been a concern as long as there have been networks.

Development of a Management Infrastructure

Until recently, only proprietary network management solutions were available. Such tools were either highly tuned for specific network-based applications or were focused on the low-level communications media (telephone circuits, modems, multiplexors, etc.) Generally, a different network management platform was required for each class of network equipment or distributed application.

Many of these management tools were of high quality and, within their limited domain, performed—and continue to perform—quite well. (Indeed, many of these specialized solutions outperform the current generation of general purpose, “open” management systems.) However, few tools were available to fill the gaps left between the vendor-specific, proprietary management offerings. Users were left to devise largely ad hoc solutions.

Over the last decade, open protocols (such as TCP/IP) and multi-vendor networks have become common. Users became free to pick and choose the best and most cost-effective equipment for each particular job. For the most part these networks performed reasonably well. It soon became apparent, however, that something had been lost. Users no longer had the same degree of management capability they had in their proprietary networks.

Management solutions for these multi-vendor networks were an afterthought. Most proprietary solutions were not designed to be extended to the heterogeneous environment. Many devices were not designed to be managed at all.

NetView

In the middle of the last decade, IBM introduced *NetView*, perhaps the first comprehensive management system, aimed at managing SNA networks. Although clearly the first great accomplishment in network management—it persists today and will continue far into the future—NetView was still a proprietary solution, one that would not serve as a model for managing anything beyond the large IBM installed base.

Similarly, the management systems developed in the telephony world, despite providing a very high level of service to users, did not serve as a complete model for managing multi-vendor networks. Layered protocol networks more closely resembled complicated distributed systems than a telephone switching system.

Toward the end of the decade, efforts began towards the definition of standards for network management. In the ISO/OSI community, an ambitious approach was taken to develop a comprehensive, all-encompassing management solution. In the TCP/IP community, the approach was more incremental, with an emphasis on experimentation and the development of progressively more powerful capabilities, each generation being based on what was learned from its predecessors.

SNMP and CMOT

While the ISO/OSI management standards were slowly evolving, the *Internet Activities Board* (IAB), after a period of debate, made recommendations for the network management standards to be used in TCP/IP networks. An extension of the *Simple Gateway Monitoring Protocol* (SGMP) which was being used by Internet Protocol (IP) router vendors was recommended as the short term solution. This was the *Simple Network Management Protocol* (SNMP). In the longer term (and as a transition strategy for OSI), the *Common Management Information Protocol (CMIP)-over-TCP/IP* (CMOT) was the recommended solution. CMOT was based upon the portion of the overall OSI management standards that had begun to gel.

The focus throughout this period was on defining the infrastructure to collect and move management data. However, relatively little progress was made on applications to reduce this raw data into useful information and then, in turn, use the information to solve real management problems.

The building blocks for network management solutions were starting to emerge. The real solutions were still a long ways off.

Deployment of the Infrastructure

Today we see the basic management infrastructure being widely deployed. A key element in this is the success of SNMP. Its popularity is enormous. At INTEROP 90, more than fifty vendors demonstrated SNMP capabilities in their products, including devices for which SNMP support is, perhaps, surprising (MAC-layer bridges, LAN hubs, and other non-TCP/IP equipment).

Many SNMP vendors

SNMP interoperability has been demonstrated for three consecutive years with a consistently growing list of vendors. Customers treat SNMP as a checklist item, and vendors have responded. Within the relatively simple framework defined by SNMP, vendors have found ways to innovate and differentiate themselves. Extensions to the standard *Management Information Base* (MIB) support management of vendor-specific features and even new standard MIBs in non-TCP/IP protocol areas (e.g., FDDI).

Meanwhile, considerable activity persists in other network management areas. Many proprietary networks and networking devices still exist, managed by often very powerful proprietary management solutions. In many cases, the proprietary solutions continue to outstrip their SNMP counterparts.

OSI Network Management

Work on the OSI management standards continues. A consortium of vendors, the *OSI Network Management Forum*, has focused on demonstrating interoperability in a subset of the overall standards. [See the article about the OSI/NM Forum in *ConneXions*, Volume 4, No. 10, October 1990.] Results of this work should begin appearing in the next year or so. The overall progress of the OSI management work, however, is slow. Hence, a transition solution such as CMOT has gained very little support.

OSI management solutions will come but are unlikely to displace the growing SNMP installed base. Real network management solutions, however, still seem to be a long way off. The focus in network management continues to be on the infrastructure, with persistent debates over protocols and related management data definitions. In the opinion of the authors, questions such as "Is CMIP better than SNMP?" are largely irrelevant to the significant issues of network management. Real network management solutions should make the underlying infrastructure transparent.

Network Management Directions (*continued*)

The central question in network management has to be: Are *real* solutions ever going to come? Network users want a network management system that plugs in, starts, and does everything. Such a system is not yet available.

Network management is in a period of experimentation. Partial solutions, such as SNMP or CMIP, provide an infrastructure for this experimentation—we can start to determine the operational measurements that really reflect the health of a network and the most effective controls to alter network operation.

The basic definitions of network management need to be reexamined. Network management cannot be contained in five magic words: *Configuration, Fault, Performance, Security, and Accounting Management*. Real network management must solve real problems faced by real people.

Requirements

The focus must shift to consider what benefits and services network management must provide to the businesses and enterprises which own and operate the networks. Network management tools must satisfy the distinct requirements of the business administrator (who views the network as an asset and a competitive tool) and the technician (who views the network as computers and communications links). A key element in all cases is reducing the overall need for human intervention.

Real network management solutions may have to:

- Consider the network as a total system rather than just its constituent parts
- Apply basic business management principles to networks—e.g., management by exception, goal setting, hierarchical management with delegation of authority, etc.

Solutions demand changes to the way we approach networking in general:

- The basic engineering of networks to be more stable
- Building management capabilities into networking components right from the start, rather than as an afterthought.

What's your role?

The networking community invites everyone to participate in the development of network management. Users ought not to be satisfied with current solutions, should not allow themselves to be dragged into infrastructure squabbles, and must continue to demand the solutions they really need.

KARL AUERBACH is Vice President and founder of Epilogue Technology in Capitola, California.

DENIS YARO is a Software Engineering Manager for Network Management at Sun Microsystems, Inc. in Mountain View, California.

INTEROP 90 Wrap-Up: A View from the Floor

by Daniel P. Dern

Record attendance

Over 20,000 people trekked through the INTEROP 90 exhibition and sessions in San Jose between October 8 and 12, 1990. That's twice the attendance of last year's INTEROP show. To provide another interesting metric, the entire ARPANET e-mail community a decade and a half ago was about this big—and today, the network user community in a large corporation or university can easily exceed this size.

And they had a lot to see! 22 two-day tutorial sessions preceding the exhibition, interspersed with open-air picnic style lunches basking in the frightening perfect San Jose weather. Over 200 vendors and even more exhibition booths, special demos, presentations, plenary addresses, short sessions...plus the usual not-to-be-missed wild (for the technology community) parties until far too late at night.

Interoperability was definitely the password. What began five years ago as a TCP/IP get-together has expanded to an event that takes in the gamut of interoperation.

SNMP

Last year saw the first SNMP demo; this year SNMP was everywhere including a large cooperative special demonstration on the show floor. The variety of application software and applications of SNMP was also much greater; this year even saw the arrival of the *Internet Toaster*, a barking toy dog, and a CD player, all controlled by SNMP.

Technology demonstrations

This year brought a wide variety of multi-vendor technology and product demos, up and down the stack: ISDN; ONC/NFS interoperation; The regional Bell Operating Companies (RBOCs) hawking SMDS, their new Switched Multimegabit Data Service; X.500 directory services (including instant digitization of your photo image); RPC (Remote Procedure Calls) in the distributed interoperable application arena; Fiber, in the form of FDDI; 10BASE-T; Plus routers, bridges, routers, analyzers, connectors and lots more.

Perhaps the most remarkable was the breadth of vendors and technologies: Sun, IBM, Digital, the RBOCs, HP, Apple, Prime...not just networking vendors, but system vendors with a growing stake in interoperable connectivity. How things have changed!

Tutorials and sessions

Many came to listen and learn as well as touch'n'test. INTEROP 90 kicked off with nearly two dozen two-day tutorials, followed by forty-five 90-minute conference sessions, plus informal Birds-Of-a-Feather (BOF) get-togethers competing with the vendor hospitality parties. Prof. Doug Comer reprised his TCP/IP overview; other topics in the curriculum ranged from X.400, X.500, the ISO Development Environment, and distributed computing applications to ISDN, GOSIP, FDDI, the X Window System, the Point-to-Point Protocol (PPP), SNMP, Network Security, Distributed File Systems, Internet Naming and Directory Services, and Gigabit network architectures—something for everybody, and nowhere near enough time to be everywhere.

Expansion

The theme of Internet expansion ran through the show, from the opening of communications with Eastern Europe and the Soviet Union, to new developments in the NSFnet, the new Canadian backbone and mid-level networks, and more commercial IP activity in the US. Alternet, announced last year by Uunet, is up and running; PSInet announced Virtual Private Network service not just for TCP/IP, OSI or DEC stacks, but for Apple, Novell and others.

continued on next page

INTEROP 90 Wrap-Up (continued)

The shownet

INTEROP 90 continued another INTEROP tradition—rolling out the Show and Tel-Net, a campus internet for use in the show. All exhibited systems capable of connecting to the shownet must, to demonstrate connectivity and interoperability; likewise, the shownet, and several banks of terminals, offered Internet members remote login and e-mail access to distant home accounts (with the usual network and server congestion periods).

This year, the network team had only eight hours available to deploy over 25 miles of cabling, so they stuck to a 10BASE-T backbone. On the whole, the INTEROP 90 shownet was a rousing success. (Especially when you consider the number of vendors, mix of equipment, and up-through-the-last-minute surprises and changes involved.)

The Internet Toaster

And hooked up to the shownet, featured in Interop Inc.'s own center floor booth, as promised, John Romkey and Epilogue Technology brought the *Internet Toaster*, running under SNMP management. (Romkey apparently has a garage full of test toast.) On Friday afternoon when Cliff Stoll (of *Cuckoo's Egg* and *NOVA* fame) strolled by, Romkey pleaded "please don't eat our demo," as Stoll helped himself to toast. Elsewhere in the international appliance networking arena, Australian Simon Hackett, working with Silicon Valley based TGV Inc. brought up their CD player under the entertainment MIB.

Quotes

Some comments from the show floor:

Vint Cerf (finally hitting the floor after four days in meetings and sessions): "All I've heard from people is how excited they are about the scale of the system, the number of network interconnected, the diversity of protocols and software that are functioning in that system."

Marshall Rose, Principal Scientist, PSI: "From exactly where we are standing, look what it says under the IBM logo in their booth... 'Because it all has to work together.' Can you imagine a statement like that from *any* vendor five years ago? The answer is No. The INTEROP show highlights how far we've come in interoperability."

(*Anon.*): "Labels such as FDDI and X are actually starting to mean interoperability and standards, not just labels that everyone wrings their own variations on."

Charles MacMullen, BICC Data Networks: "It's exciting how the vendors come together at INTEROP, work together to make something like promoting 10BASE-T happen, instead of simply being competitors. It's important to end users for us to demonstrate multi-vendor interoperability—and this is the forum to do it in."

Jim Miles, Information Week: "INTEROP cuts across technologies and vendors, and sheds light on important technologies that might otherwise not get appropriate public attention. And it's one of the best organized shows I've been to."

Cliff Stoll: "It's terrific! Amazing! Connectivity! There's eagerness all over the place! It's a spider-web connecting one booth to another!"

Valery Udalov, Latvian Academy of Sciences: "It's amazing to me...to see how many products and how much live connectivity is here. It's very interesting, I've never seen so many computers in one place. In this hall are probably more computers than in all my network."

Einar Stefferud, Network Management Associates: "My reaction is the same as last year: Gee, Toto, I don't think we're in academia anymore."

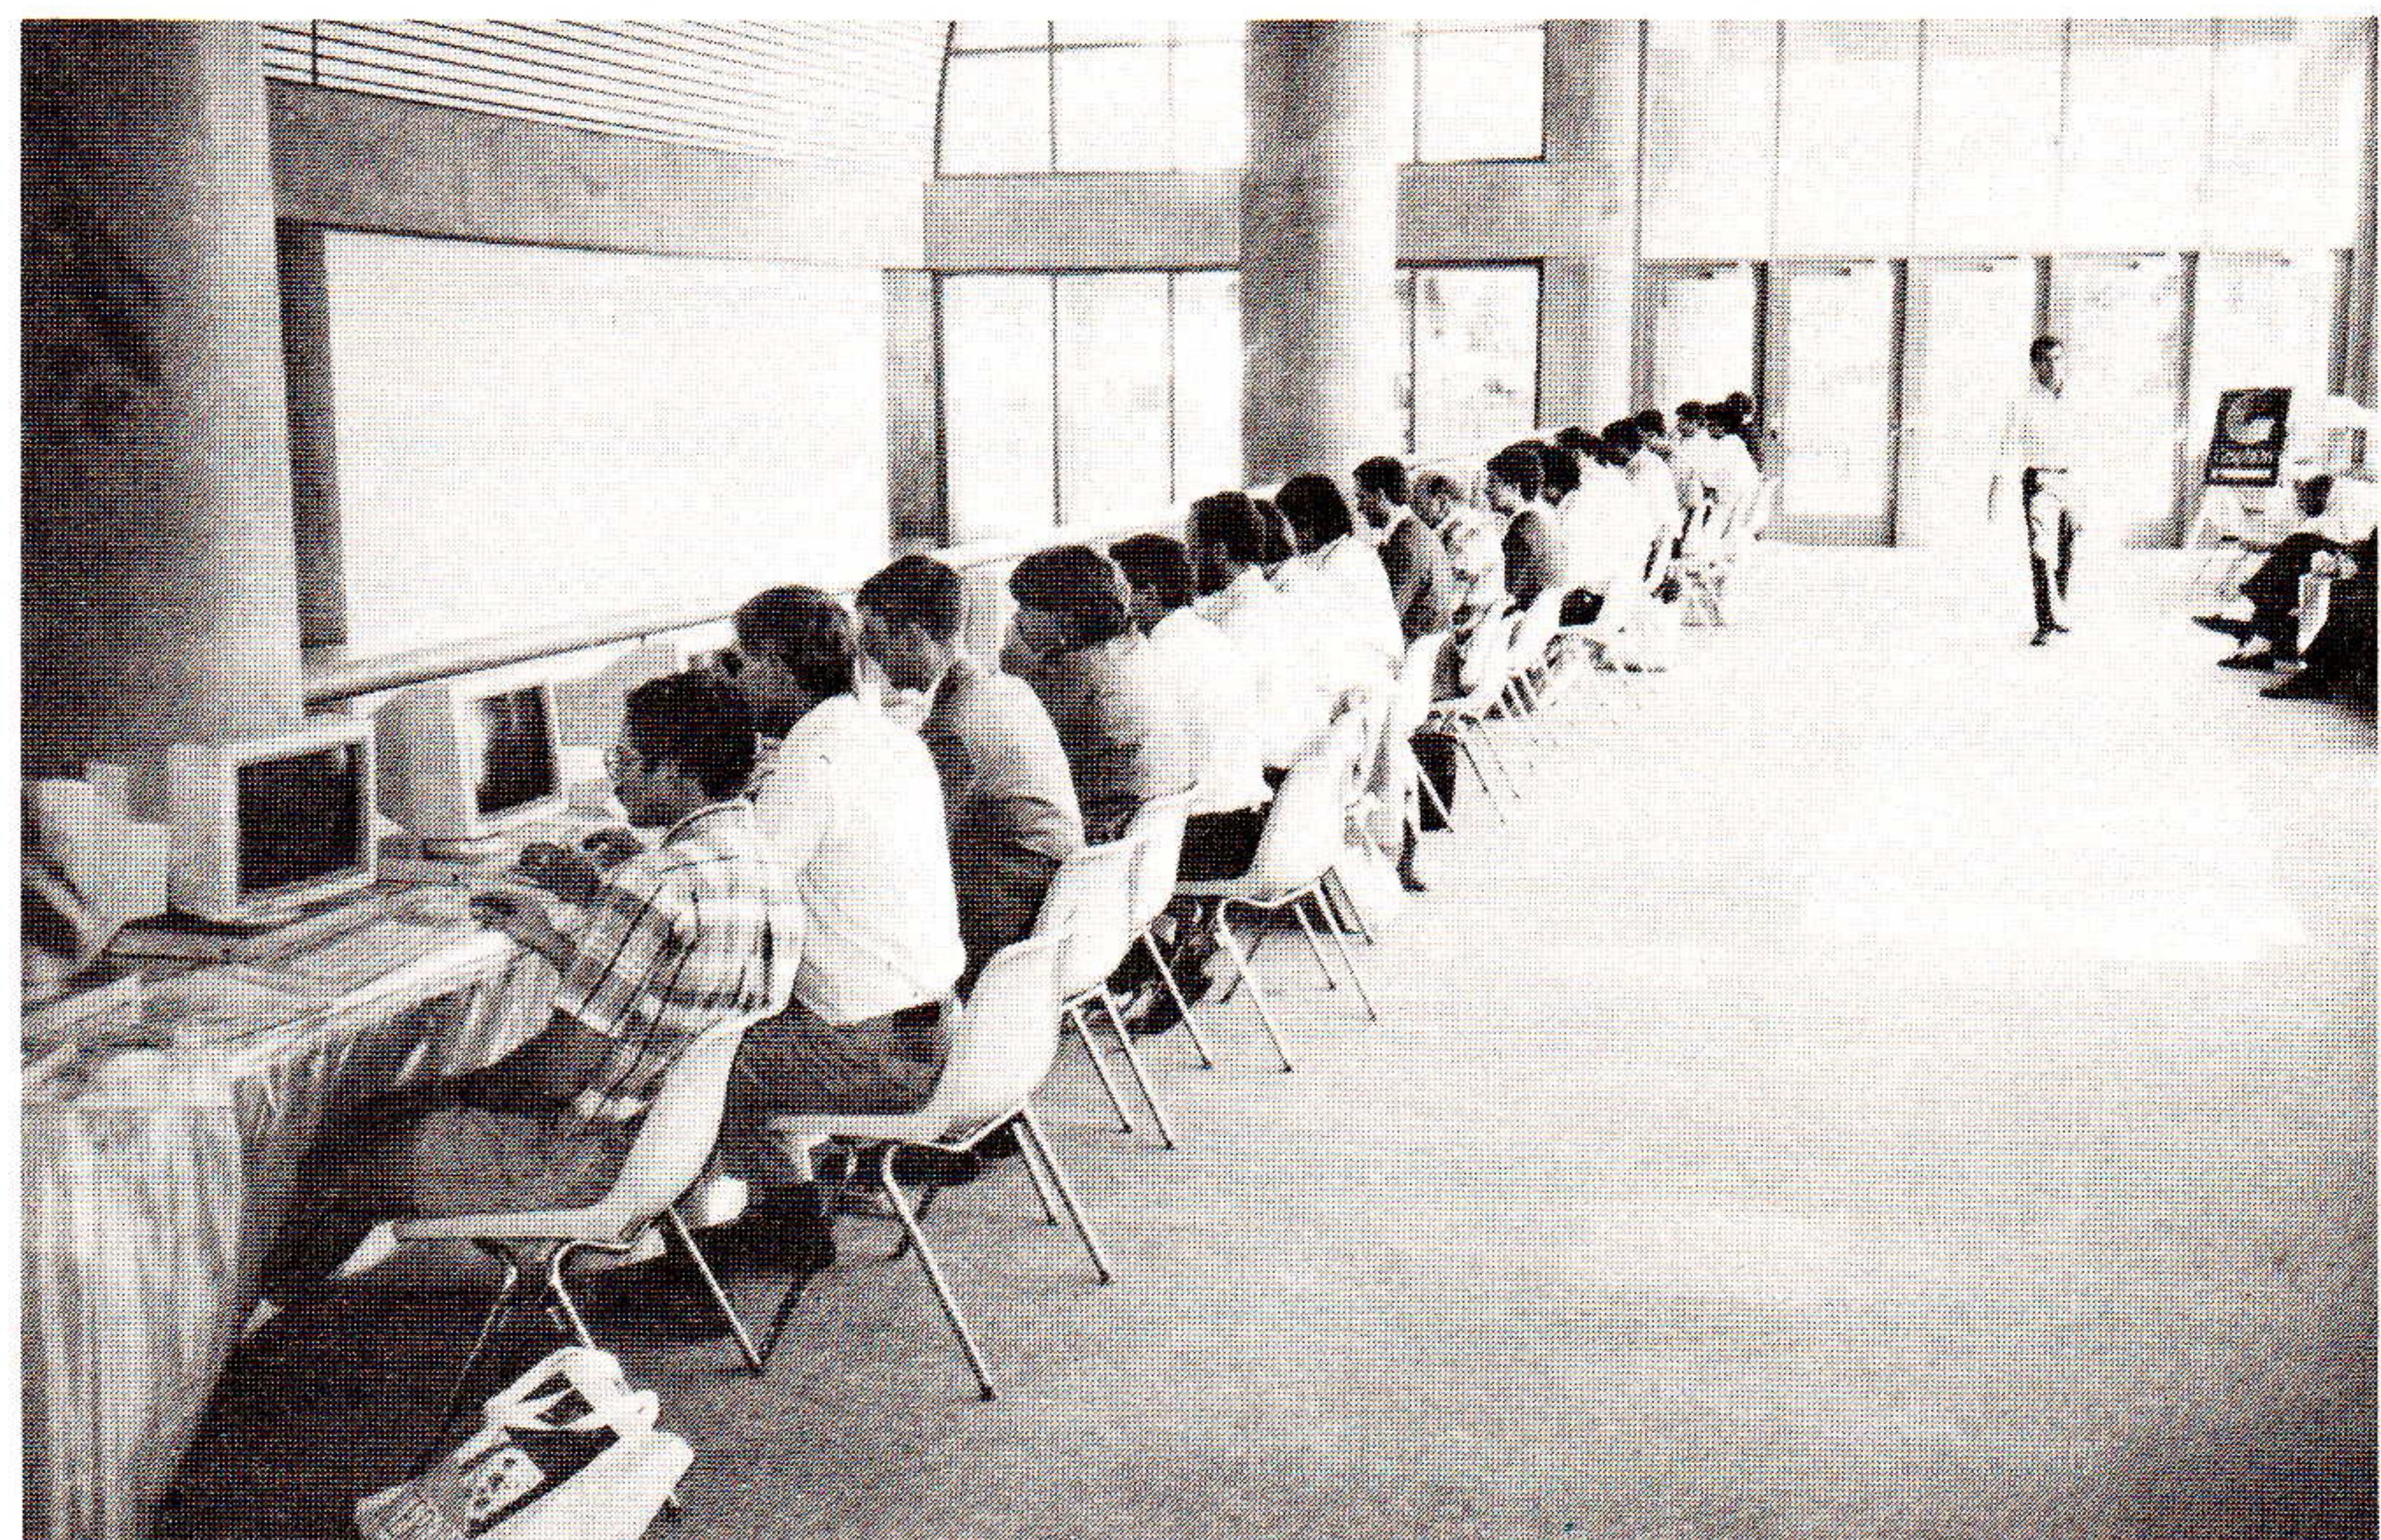
Bob Hinden, BBN: "It's very impressive how much growth there's been since I started coming to these shows. This technology has been commercialized, is being used to solve real problems for real people."

Dave Crocker, DEC and IETF: "I think it's wonderful to see the expansion of this technology into multiple businesses and multiple technologies. So that it is no longer a few people making money, it's an industry. And this event captures the fact in one place, which is just awesome."

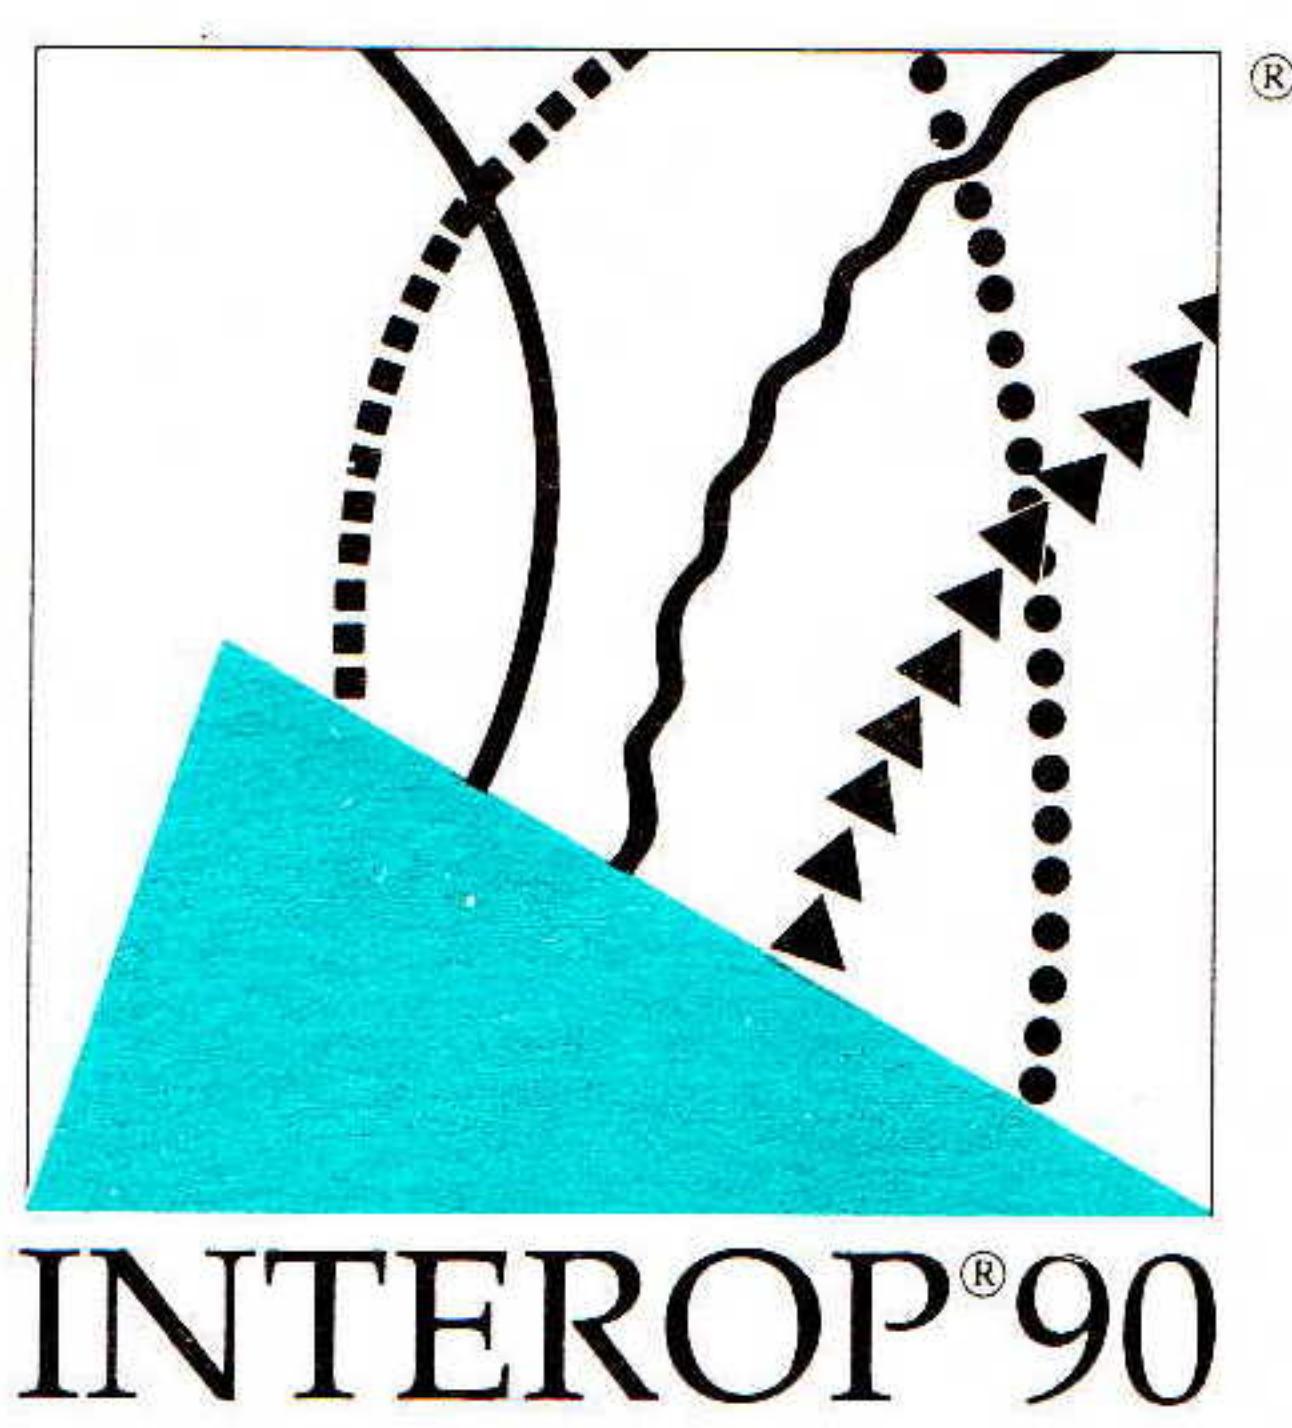
Last year seemed to mark INTEROP's crossover from a purely technical event to also a commercial, tradeshow event. INTEROP 90 confirms this is a trend, not a blip, and that interoperability starts with TCP/IP and networking, but doesn't stop there. The challenge for INTEROP 91: Successfully managing success, and guiding growth. One thing you can be sure of: like INTEROP 90, INTEROP 91 will be the place to be!

DANIEL P. DERN is a Watertown, Mass-based free-lance writer specializing in technology, science and industry. A frequent contributor to *ConneXions*, including last year's ARPANET historical retrospective, Dern writes for leading publications and vendors in the network and computer industry, as well as writing hi-tech humor columns, science fiction, musical theater, and teaching UNIX. He was previously PR Manager at BBN Communications. He can be reached via e-mail as ddern@world.std.com (Internet) or dandern (MCIMail).

Editors note: This is only a quick recap of INTEROP 90. We plan to cover the conference in more detail in subsequent issues. In particular, we will take another look behind the scenes and find out how the shownet was designed and built in record time. Technological advances and serious time constraints resulted in many design changes from the 1989 network. Now turn the page for some snapshots from INTEROP 90.

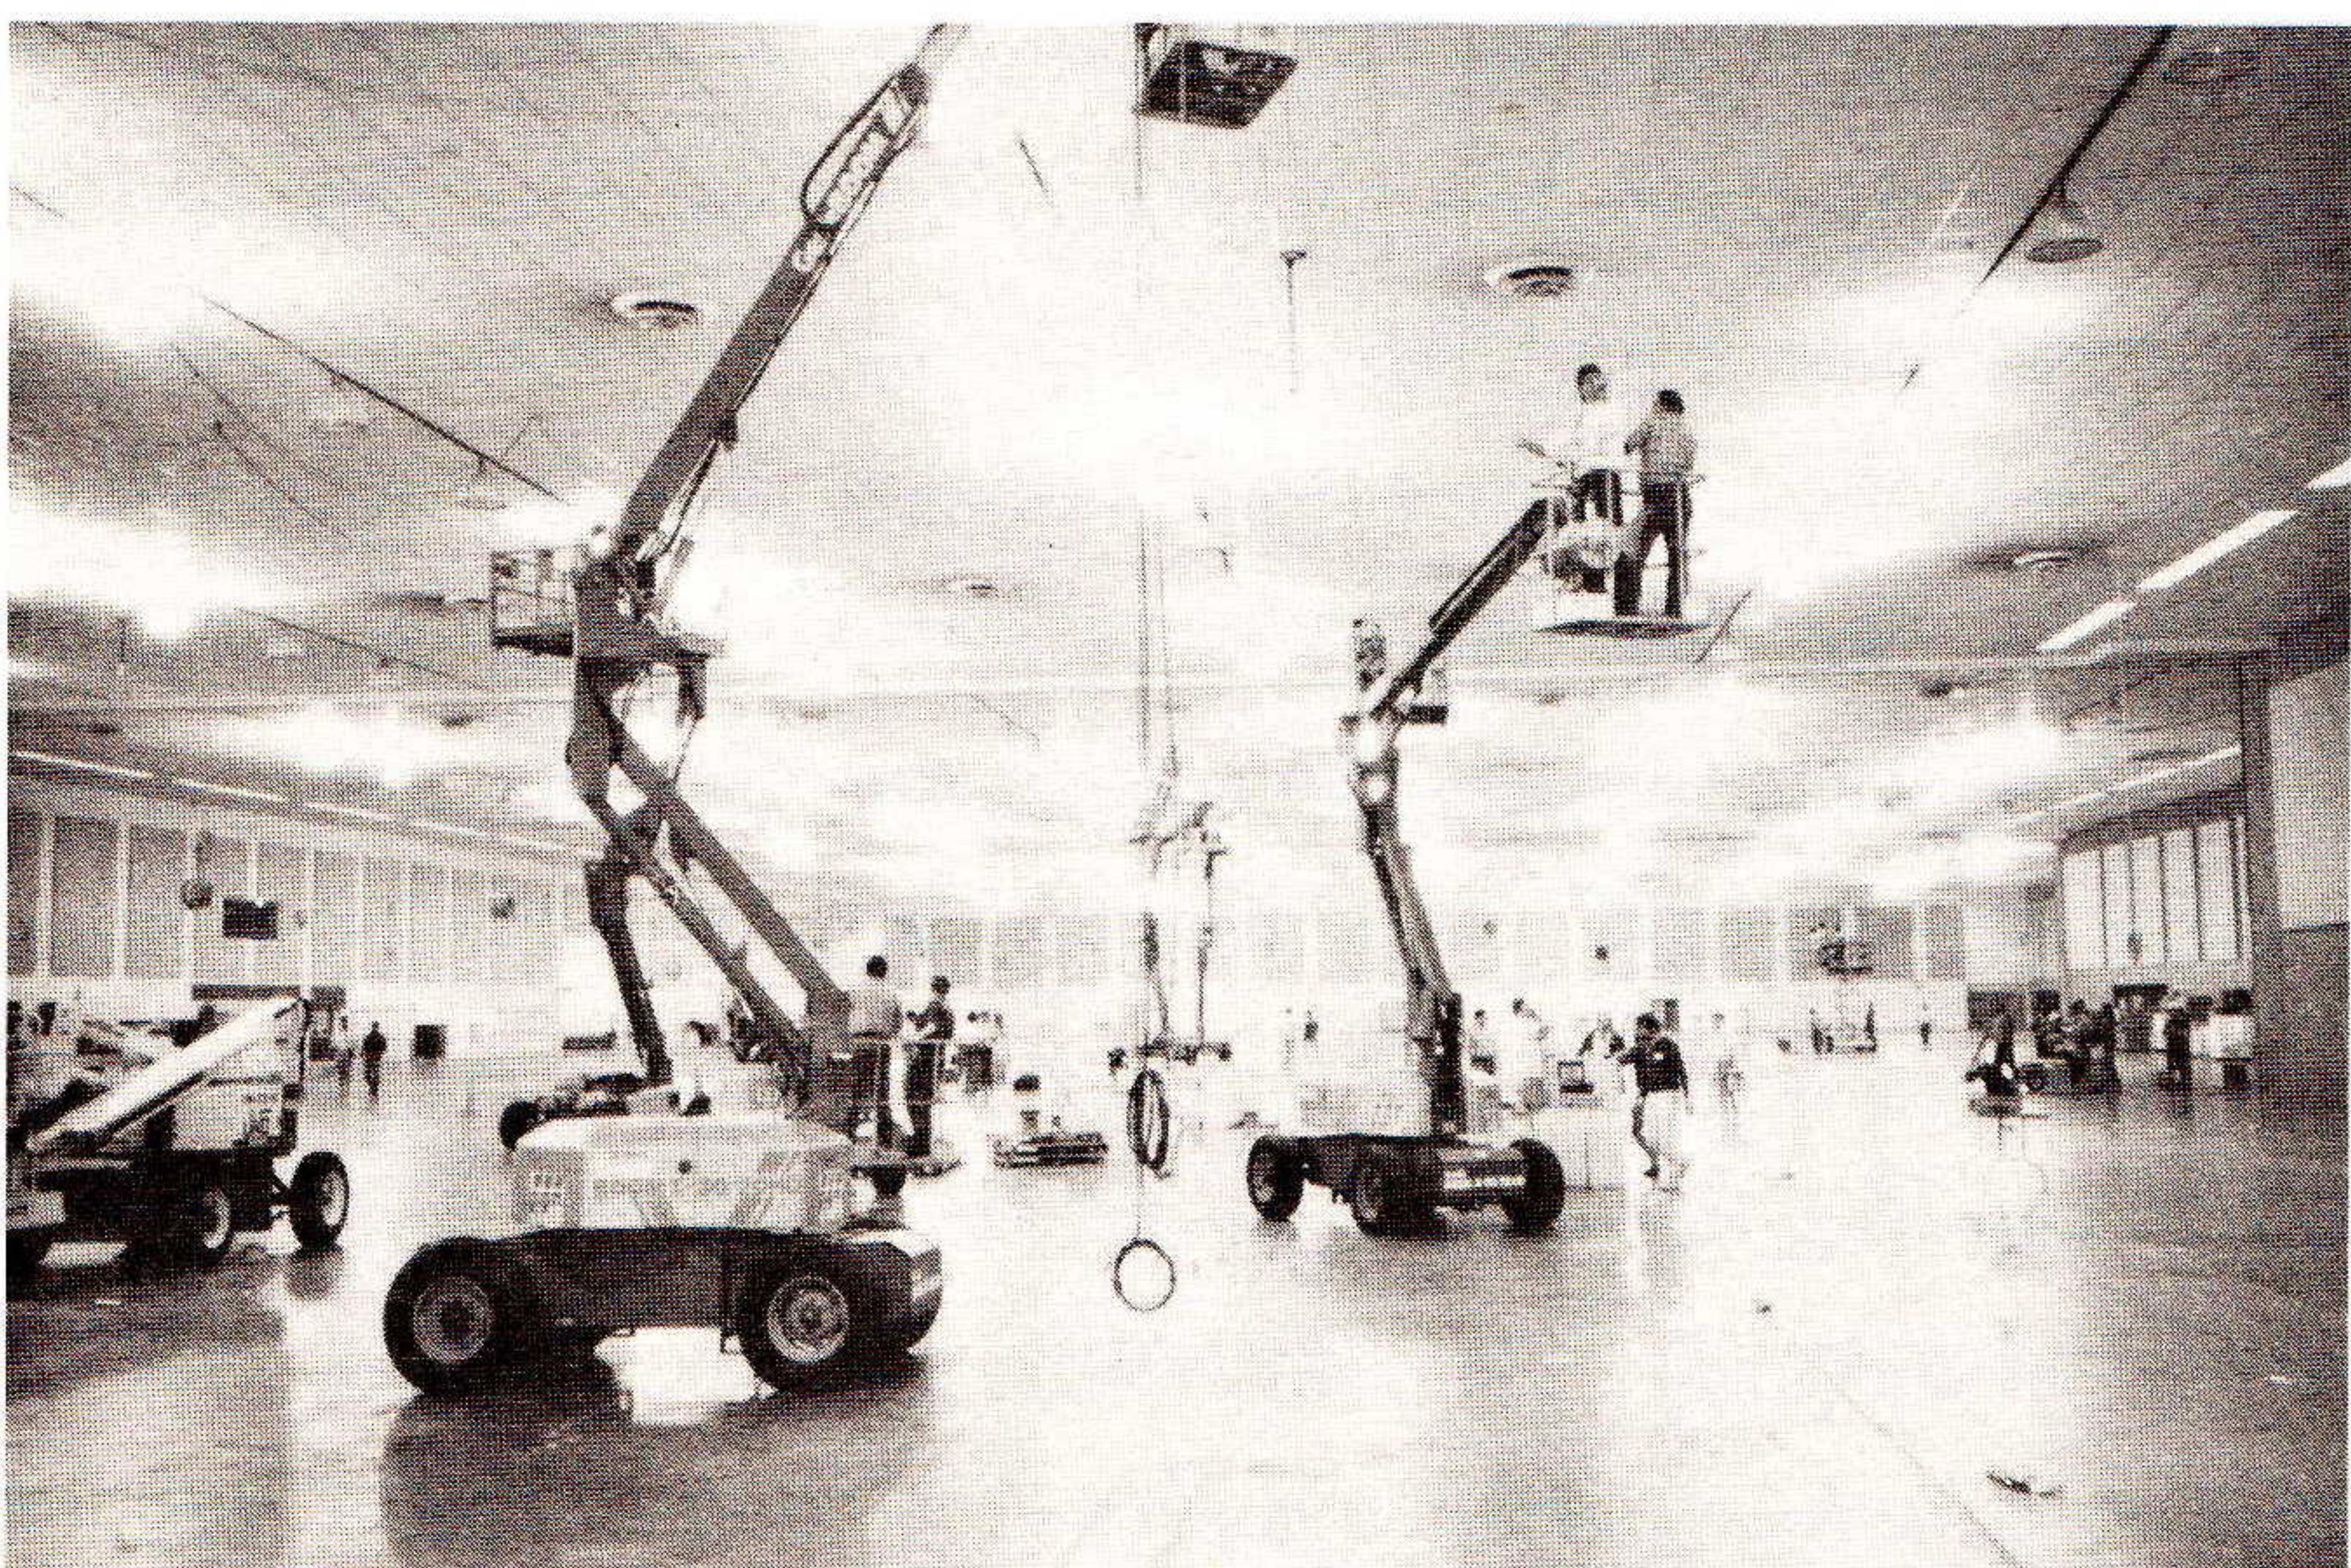
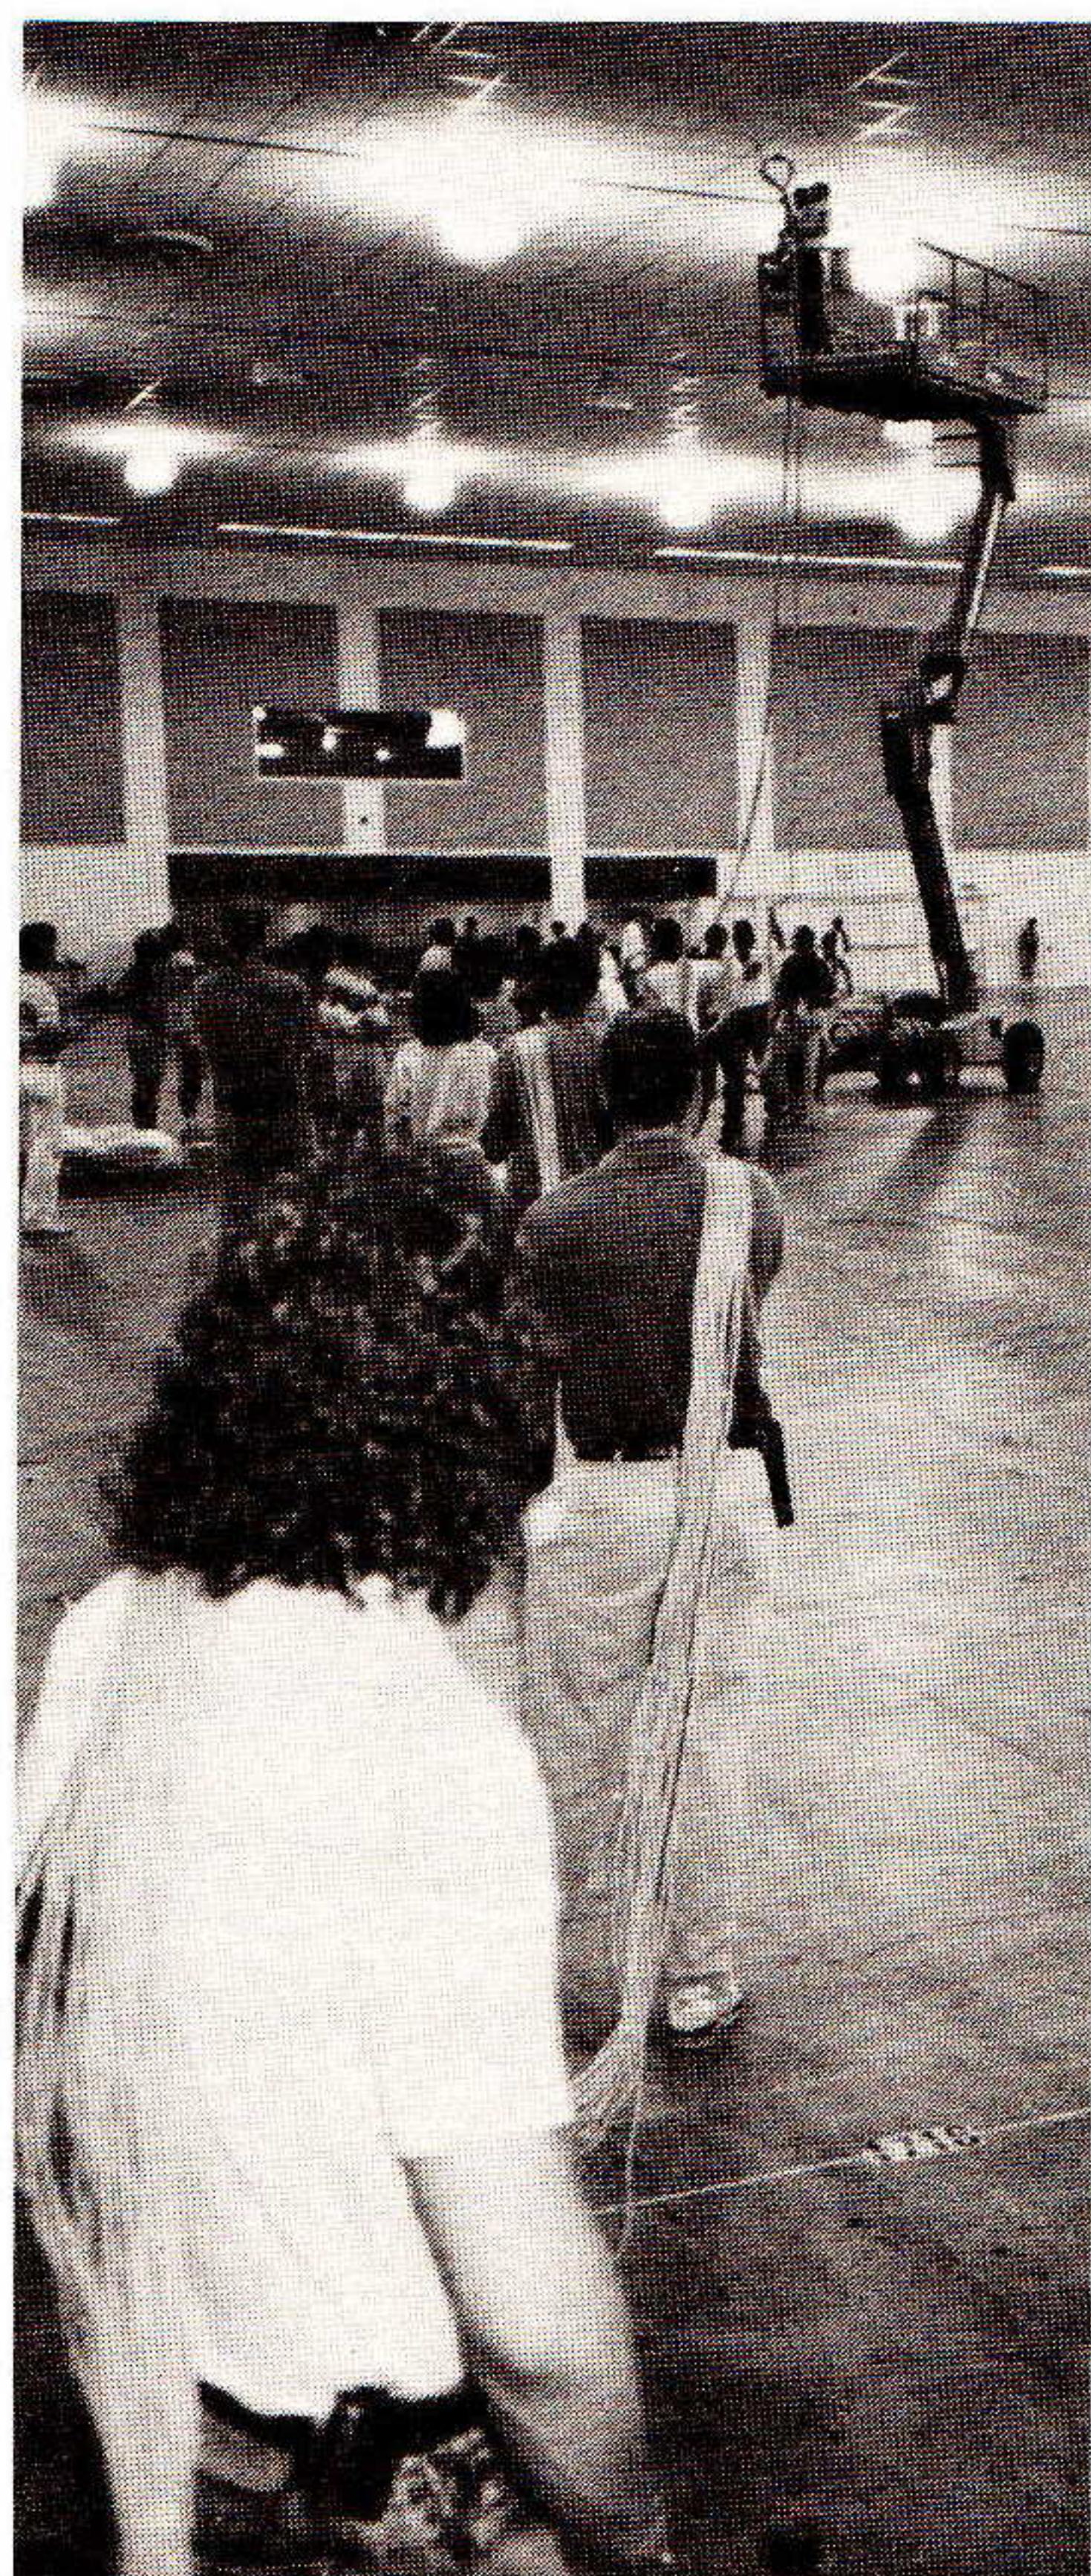


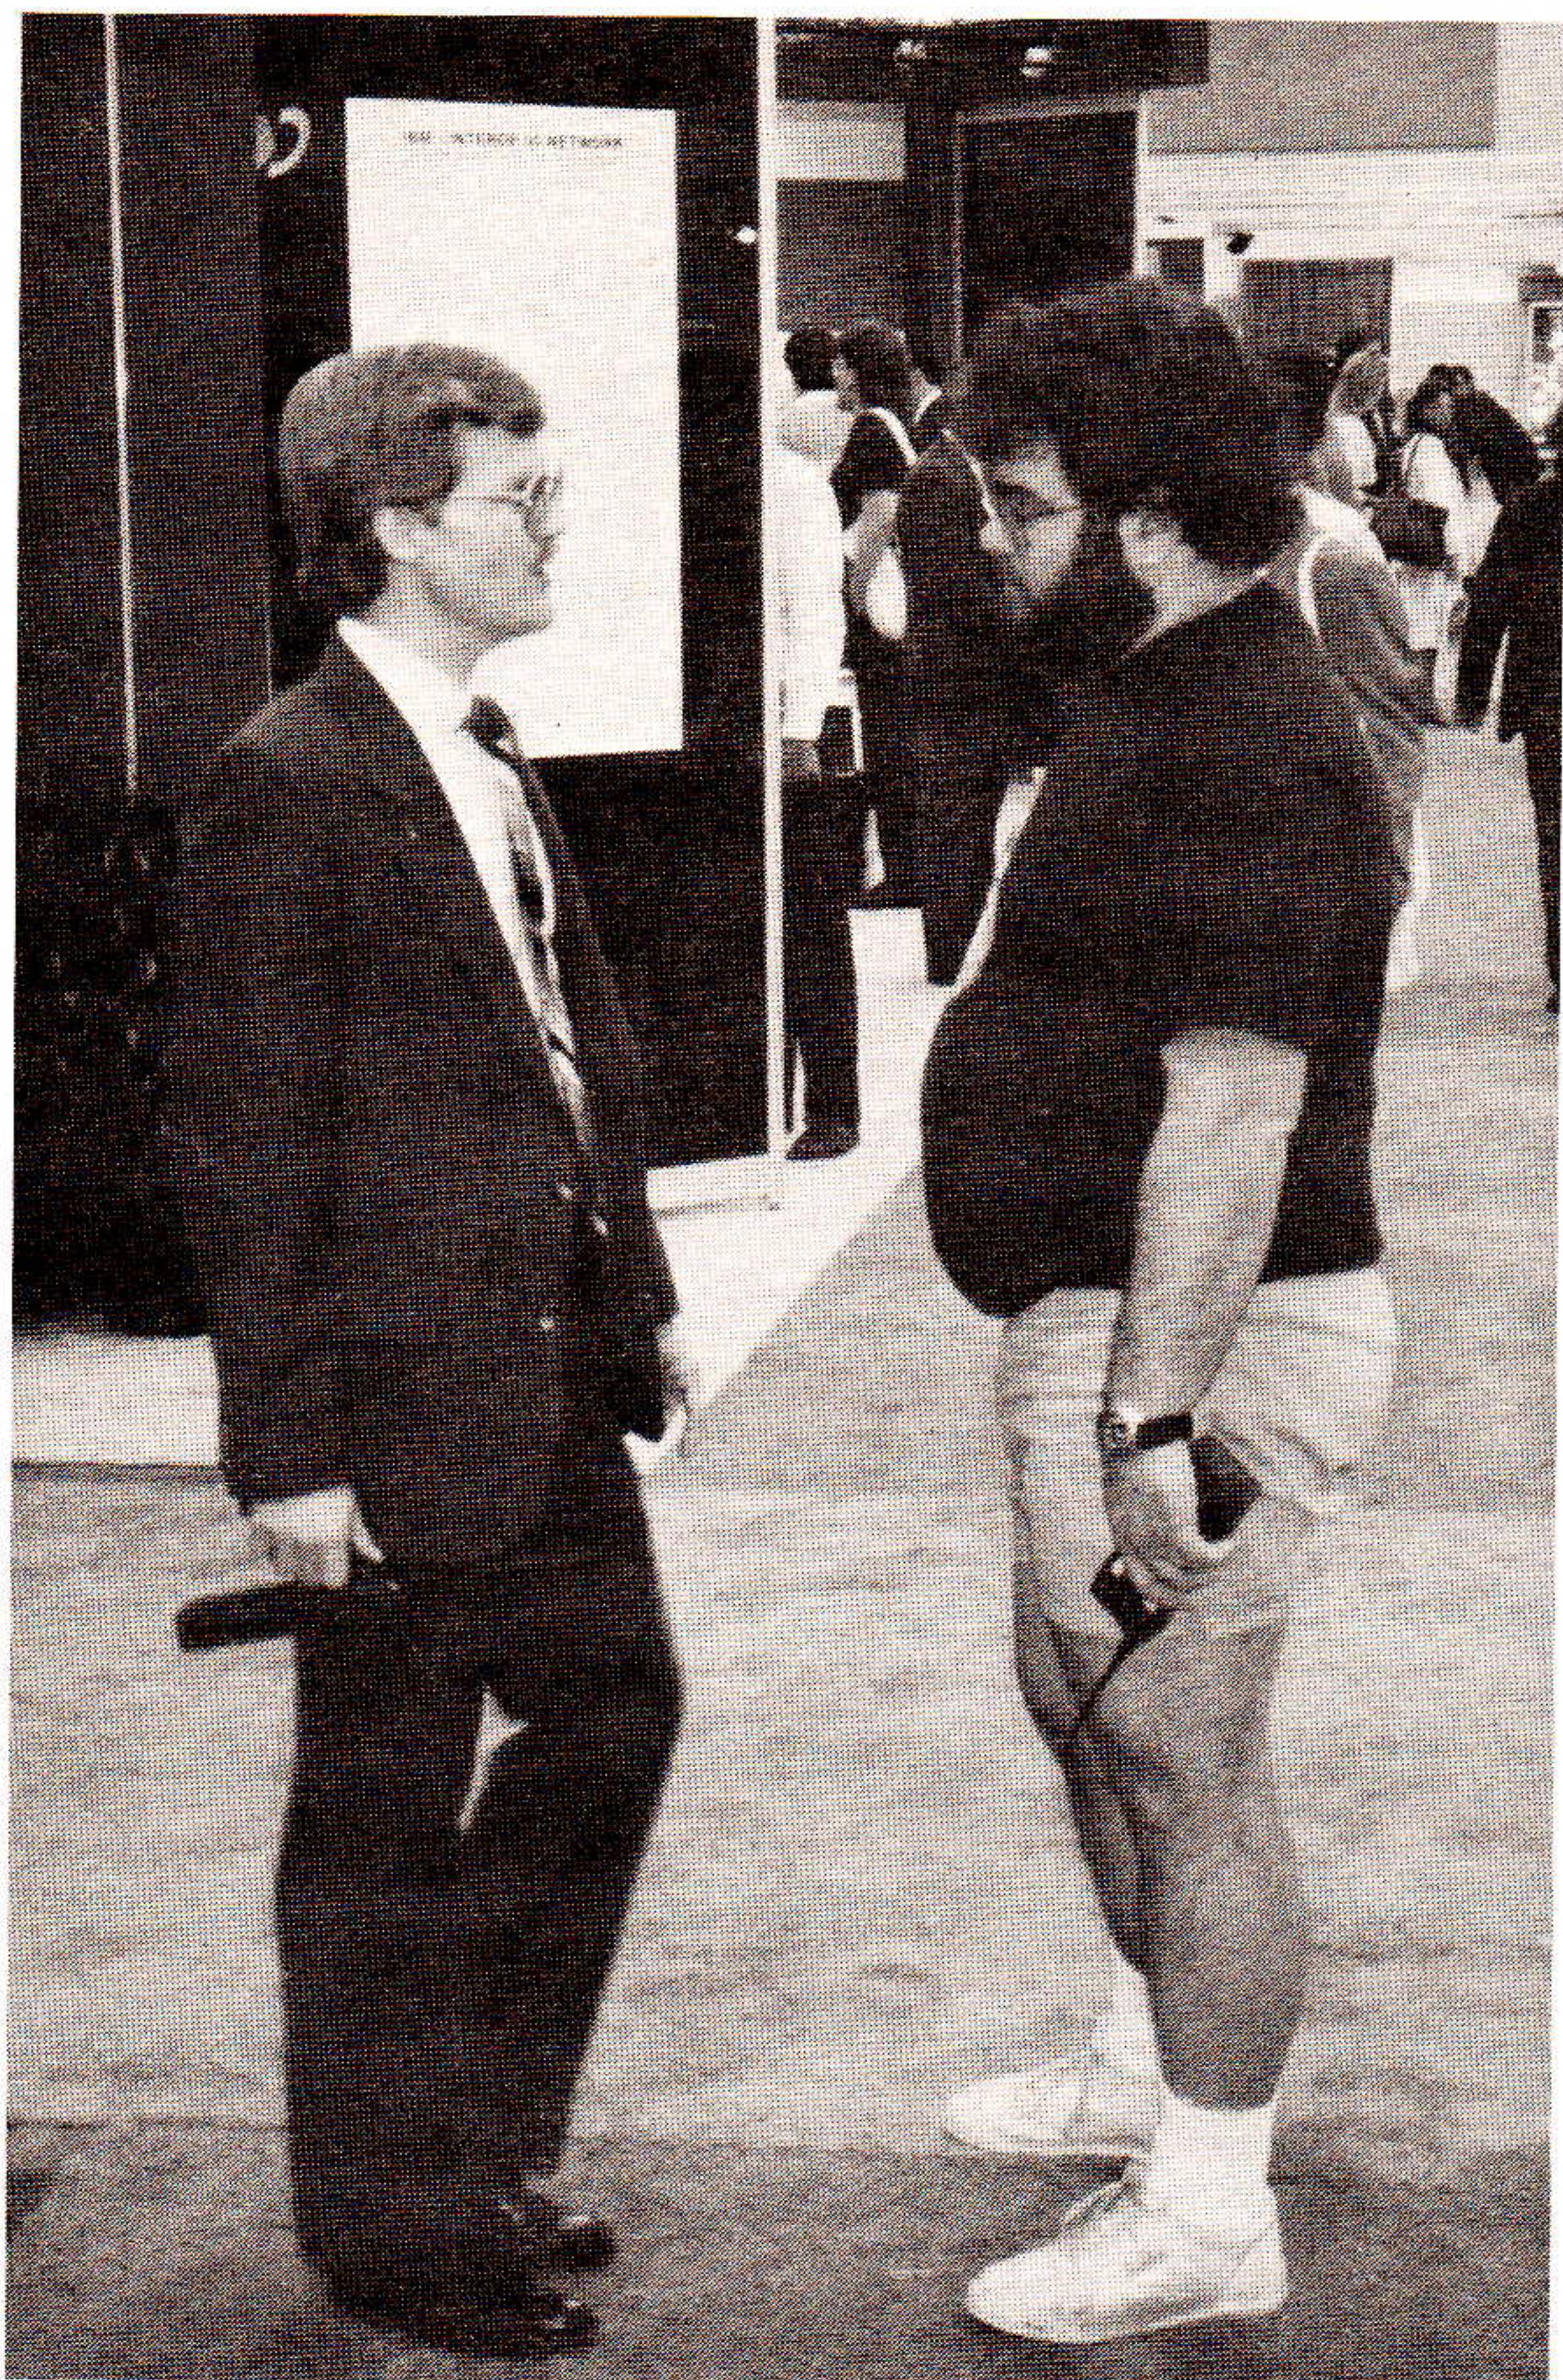
"Just checking my mail." One of the 3 INTEROP e-mail centers.



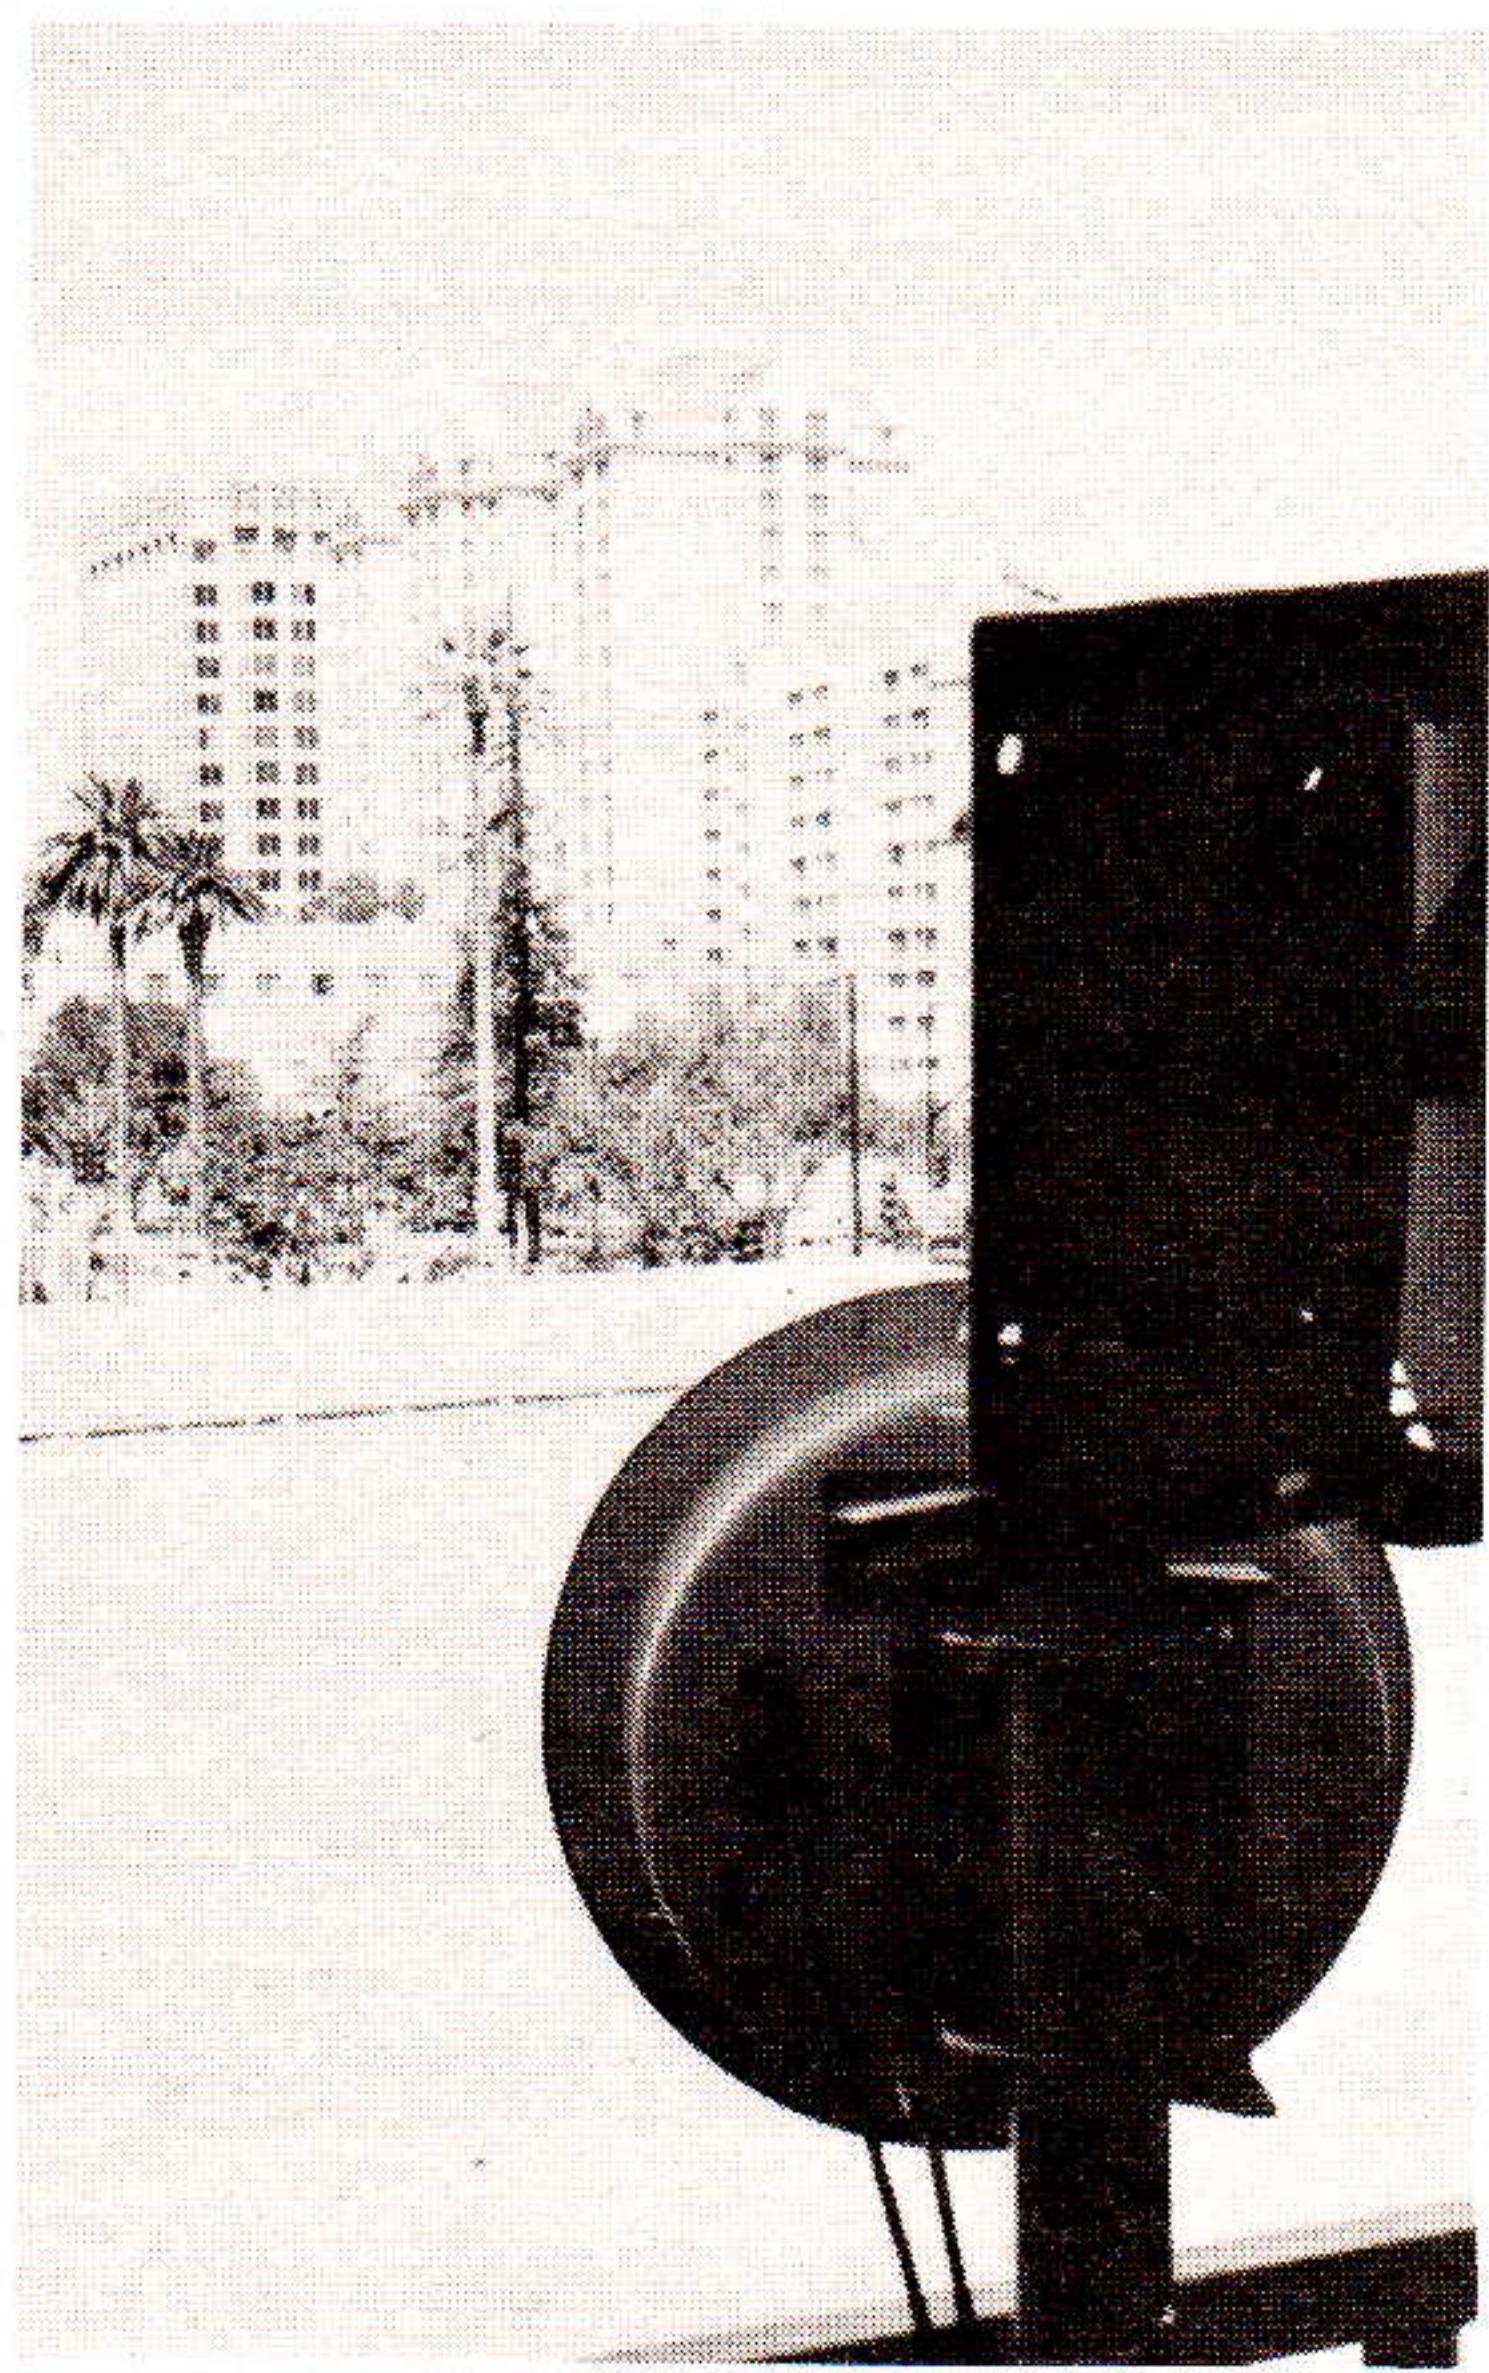
Snapshots

Late Saturday night: the network team arrives to an almost empty exhibit hall, still being cleaned from the last tradeshow. The cable for the network backbone is carefully unrolled from pallets and teams of volunteers help hoist it into position high above the floor. In eight hours hundreds of crates of vendor exhibit materials will start arriving making any movement of the "cherry pickers" and scissor lifts impossible, so the job must be done without any delay.





By Wednesday morning everything is in place ready for the more than 20,000 attendees. **Above left:** Peter de Vries, Shownet Director confers with Steve Knowles, Chief Network Engineer. Walkie-talkies let network team members communicate between different parts of the hall, and the Network Operations Center (NOC) "Eye in the Sky." **Above right:** Interop Inc's center booth includes a large array of network equipment from various vendors. Also featured is the Internet Toaster developed by John Romkey of Epilogue Technology. **Below left:** As was the case at INTEROP 89, a microwave link connected the shownet to an e-mail center at the Fairmont Hotel. **Below right:** The finished product. This photo was taken from the same spot as the one in the top left corner of page 16.



Profile: NORDUnet

by Mats Brunell, Swedish Institute for Computer Science

The program

NORDUNET (note capitalization) is a networking *program* in the Nordic countries, funded by the Nordic Council of Ministers. The program runs from 1986 to 1991. The NORDUNET program operates the NORDUnet *network*. The participating countries/networks in the NORDUNET program are:

Denmark/DENet
Finland/FUNET
Iceland/SURIS
Norway/UNINETT
Sweden/SUNET

Goals

The goal of NORDUNET is to provide harmonized network services to the Nordic research and development users in cooperation with the national research networks mentioned above. A secondary goal is to establish good inter-Nordic relations in networking. Knowledge in networking and good international contacts are vital in this respect.

History

The NORDUnet idea was born in September 1987. The so-called X.EARN project brought forward a plan containing all the pieces needed to implement the proposal. It contained service requirements, a technical solution, and estimated costs for implementation and operation. It also gave a overview of the proposed organisation.

The decision was made by six parties. The five national networks had to take full responsibility for the long-term operational costs from the beginning. The funds made available by the Nordic Council of Ministers made the decision somewhat easier. The estimated total cost for the implementation of NORDUnet is \$800,000. The operational costs estimated for 1990 is \$750,000.

The network was officially opened in October 1989. The main reason for the good results was true Nordic cooperation, and a lot of additional help from others; HEPnet, EUnet, EARN and DEC. The current backbone has an estimated life-time of 3–5 years.

Activities

The NORDUnet activities focus on provision of services. This means that we are trying to extend the services and interconnectivity to new networks to the benefit of our users. We are also planning for an introduction of OSI-based services through pilots and experiments for the future. These include X.500 Directory pilots, the harmonization of e-mail addresses and development of national e-mail gateways

The NORDUnet network is an operational service. NORDUnet has service agreements with most of the existing international network services like EARN, EUnet, NSFNET, HEPNET and SPAN.

NORDUNET takes an active part in the RARE (*Réseaux Associés pour la Recherche Européenne*) work and supports the goals of COSINE (*Cooperation for Open Systems Interconnection Networking in Europe*).

Organization

The NORDUNET steering committee reports directly to the Nordic Council of Ministers. The project team is small, only 1 part-time person is on direct contract for the coordination and development of the program.

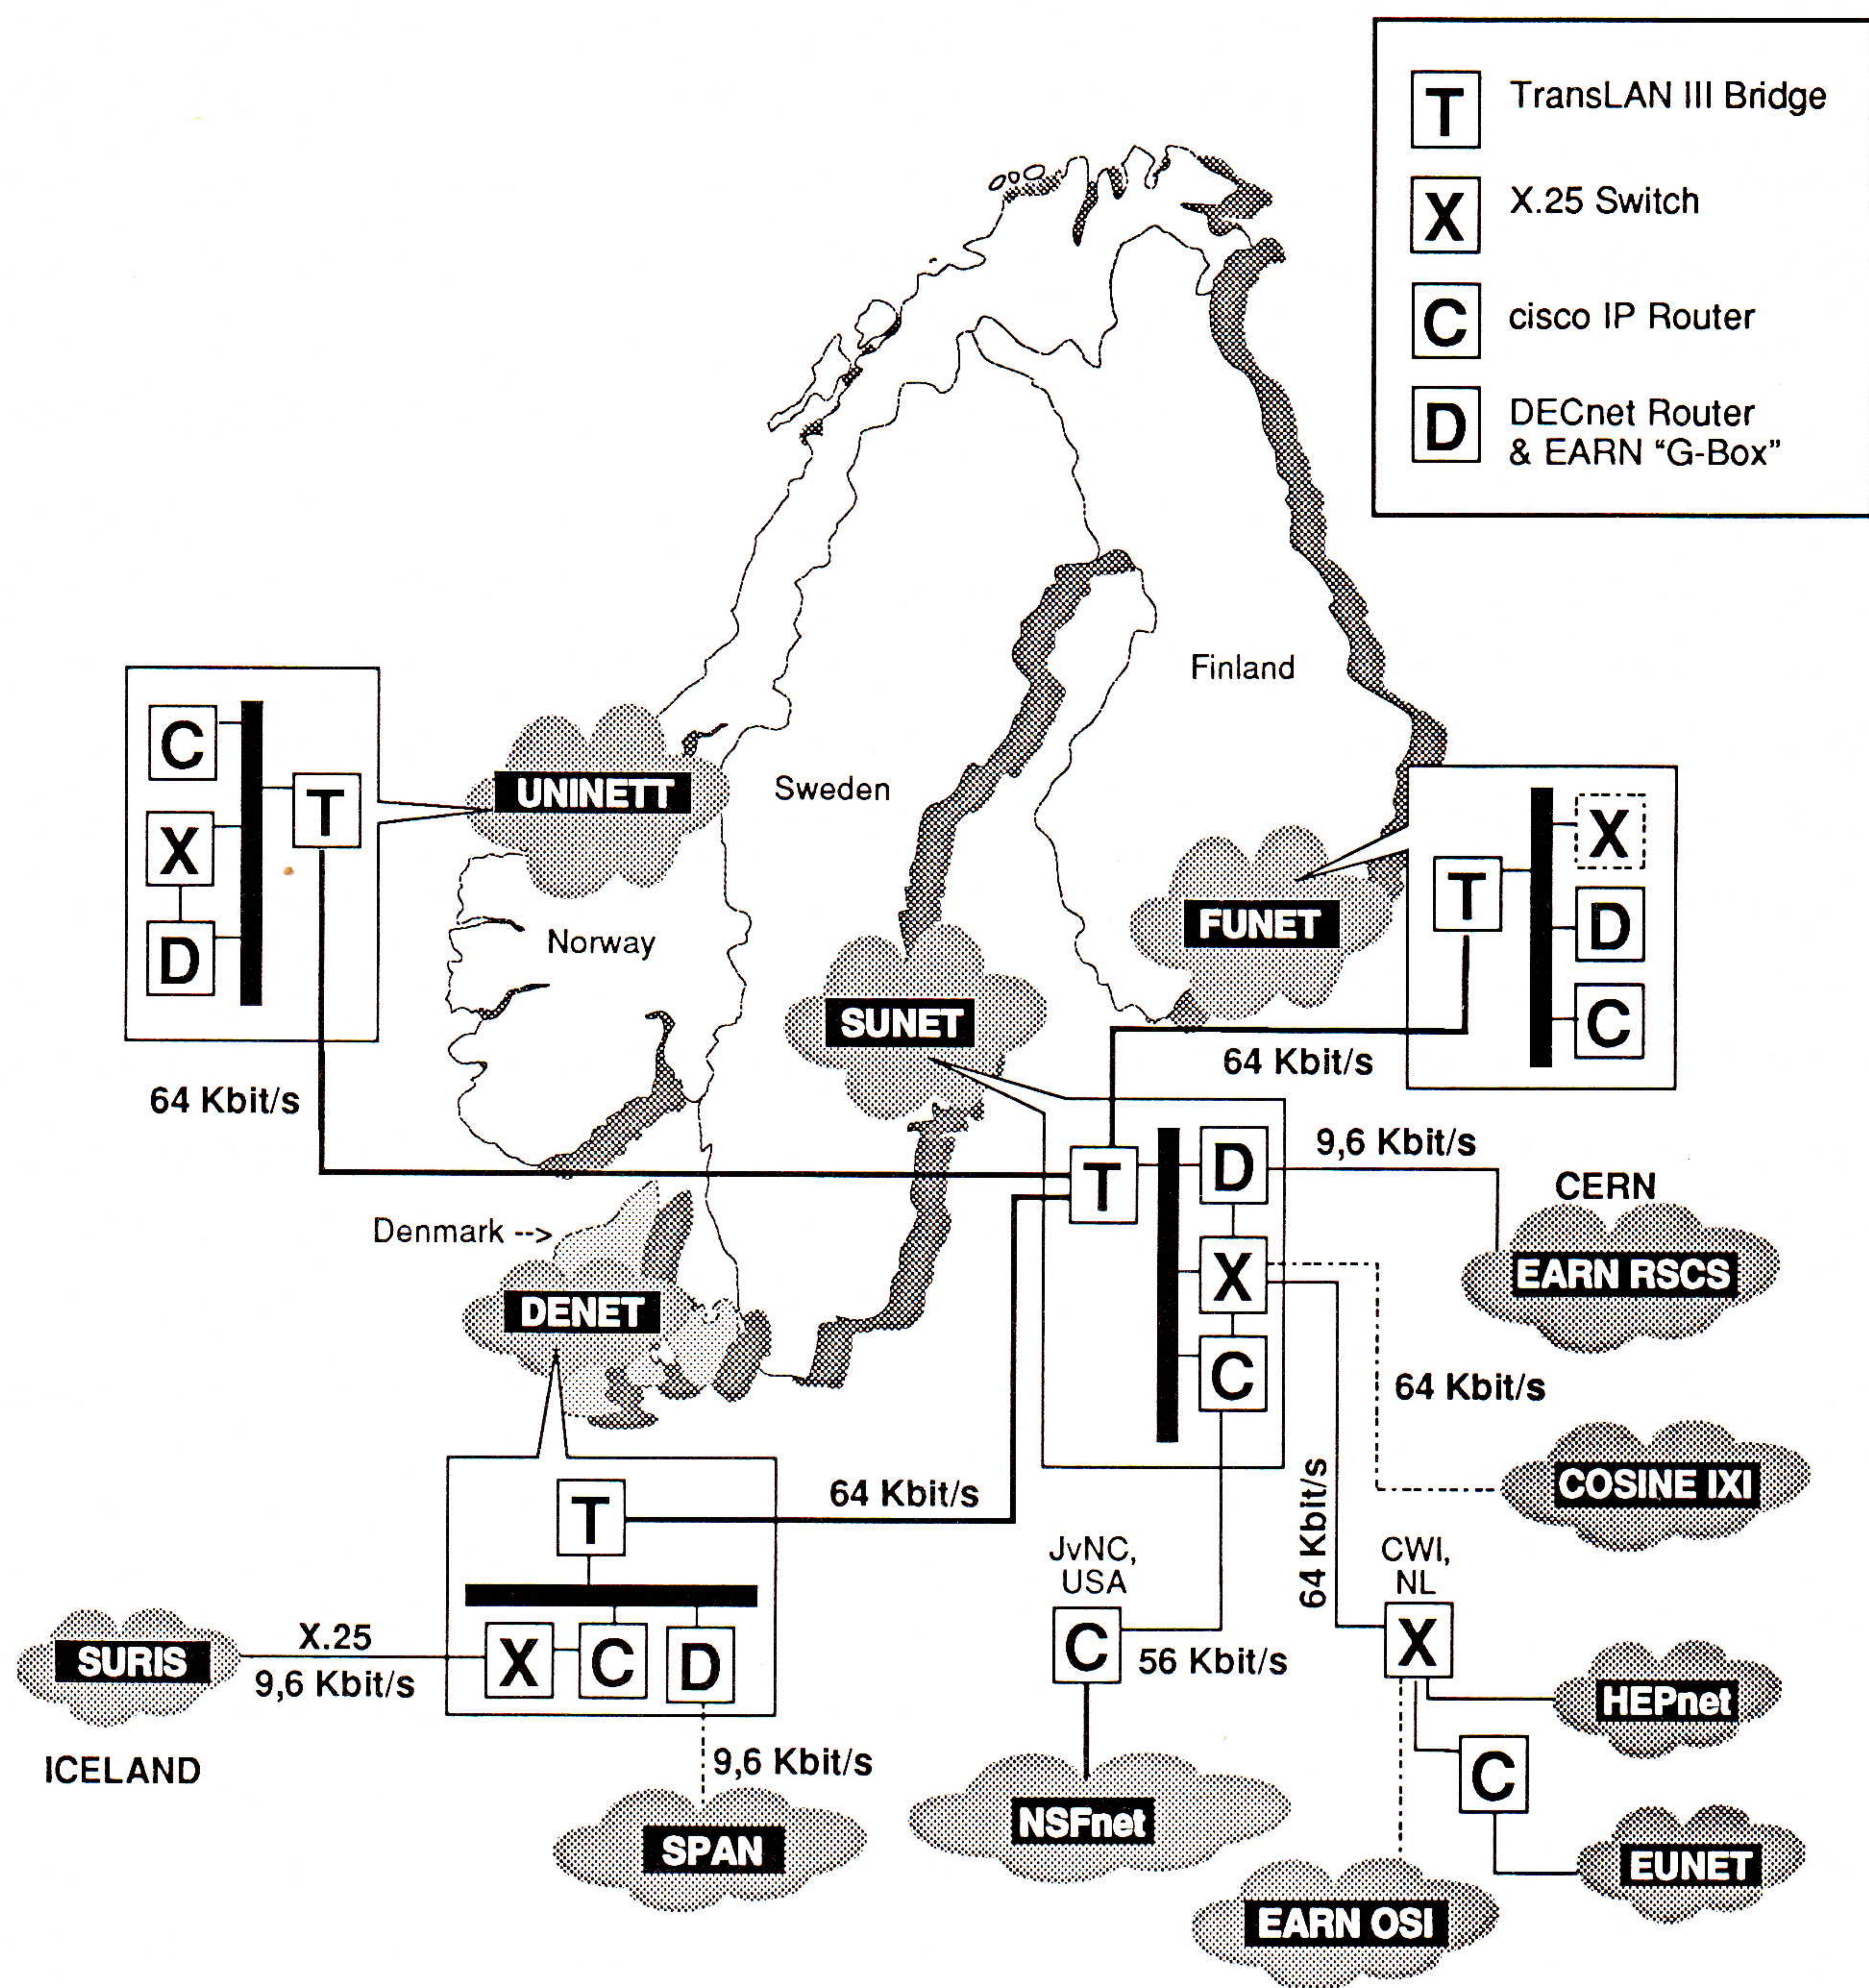


Figure 1: NORDUnet network topology

The program is hosted at the Swedish Institute of Computer Science, SICS. The work is carried out in projects under contract or in working groups. Projects like the NORDUnet implementation are normally carried out under contracts to persons and organisations taking part in the national networking activities.

The NORDUnet board has the overall responsibility for the NORDUnet network. The operations group consists of the service coordinators for the services provided. The NORDUnet staff consists of: Peter Villemoes, Chairman, UNI-C Copenhagen, Denmark, Bjorn Eriksen, Network Manager, KTH, Stockholm, Sweden and Mats Brunell, Executive Officer, SICS, Kista, Sweden.

Profile: NORDUnet (*continued*)

The national networking activities are very important in the NORDUNET work. If an implementation of the result of a plan or study is to take place, it is normally up to the national network organisation to do it.

The NORDUnet network

The NORDUnet transport network is a wide area network based on MAC-level bridges and “network-level” routers. They form a logical Ethernet connected through 64Kbps leased lines provided from Swedish Telecomm International (STI) and Scandinavian Telecommunications Services AB (STS).

NORDUnet provides, through its interconnections to the US and central Europe, access to the following networks: The Internet, BITNET/CREN, EUnet, EARN, HEPnet, SPAN and the COSINE/RARE *IXI* pilot service

Protocols and standards used

The basic network is an IEEE 802.3 Ethernet. On top of Ethernet the following protocols are provided as operational services: ARPA IP, DECnet, EARN NJE protocols, and X.25. NORDUnet will also carry the EUnet NEWS and mail on the backbone, and aim at a service agreement with EUnet in the future.

The present US connection uses an INTELSAT 56Kbps IP connection between The Royal Technical Institute (KTH) in Stockholm, and the John von Neumann National Supercomputer Center (JvNC) in Princeton, New Jersey. JvNC is one of the NSFNET backbone-sites.

A backup agreement with EUnet/UUNET gives both NORDUnet/Internet and EUnet/UUNET users a better service level. This is made possible by the square connections: KTH to JvNC, KTH to CWI, CWI to UUNET and NSFNET between UUNET/SURAnet and JvNCnet. (See Figure 1).

Iceland is currently connected to the NORDUnet by public X.25 (over a satellite link) to the cisco gateway at UNI-C/Copenhagen, Denmark.

Subnet infrastructure

The following equipment is placed on the inter-Nordic Ethernet:

- Vitalink TransLAN III, MAC-level bridges
- cisco Internet IP-routers
- DEC VAX 3600 systems for DECnet and EARN
- X.25 switches
- The NORDUnet “NIC” host nic.nordu.net
- Nameservers and network management systems

The management and operations are provided through contracts with several sites in the Nordic countries:

- Basic LAN: KTH Stockholm,
- Internet IP and the US connection: KTH Stockholm,
- DECnet, HEPNET/SPAN coordination: UNI-C Copenhagen
- X.25 and related problems: RUNIT, Trondheim
- EARN RSCS and related services: FUNET, Espoo

DECnet	<p>The national DECnets have harmonized the area numbering according to the SPAN/HEPNET scheme. This scheme uses a “Poor-Mans-Routing” and a <i>max area</i> gateway to the international HEPnet and SPAN DECnet backbone, where the Nordic countries have one <i>area number</i> under the <i>max number</i> 46 (area 21).</p> <p>This means that the Nordic DECnets use the area numbers between 47–63. The area number 21 is allocated for connection to the international DECnet backbone. The following distribution is used:</p> <table> <tr> <td>Finland:</td><td>47–49, 61–62</td></tr> <tr> <td>Denmark:</td><td>50–53</td></tr> <tr> <td>Norway:</td><td>54–55</td></tr> <tr> <td>Sweden:</td><td>56–60</td></tr> </table>	Finland:	47–49, 61–62	Denmark:	50–53	Norway:	54–55	Sweden:	56–60
Finland:	47–49, 61–62								
Denmark:	50–53								
Norway:	54–55								
Sweden:	56–60								
Internet IP	<p>The Internet-IP is one common routing domain. The routers internally uses IGRP and RIP. Externally EGP is used. The primary route to the Internet is via JvNC. More than 20 class B networks have “connected” status. They are being announced at relevant gateways. The backup link from CWI to UUNET uses the cisco “administrative distance” function.</p>								
X.25	<p>X.25-based service is provided on top of the Ethernet. The addressing problem is, however, complicated. The existing X.25 services are based on the CCITT X.121 addressing scheme. HEPNET/X.25 and others have similar schemes. NORDUNET has made a X.121 address plan for use with the existing X.25 networks nationally and internationally. A lot of address mapping has to be applied.</p>								
Lessons learned	<p>Some of the experiences gained during the implementation of the NORDUnet network are summarised below.</p>								
The PTT situation	<p>STS/STI installed the 64Kbps links with an average installation time of 3 months. A private muxed network of 2Mbps or higher speeds provide Point Of Presence (POP) connections in the main capitals. From there, local loops are used from the national or local PTTs.</p> <p>A line quality problem and bad availability has been degrading our service. The availability on the worst links is only 96–97%, the best being the SE to NL one with up to 99%. The bit error rates can be as high as 1 in 100,000 bits.</p>								
Bridge equipment	<p>The Vitalink TransLAN III equipment used in NORDUnet has been very reliable. The only problem has been mixing connection-oriented and connectionless traffic to a loaded bridge/link. This is a general problem, not Vitalink’s in particular.</p>								
IP implementation	<p>NORDUnet uses cisco IP routers. They have proven to be reliable once initial faults had been “burned out.” The IP service has proven to work very well and was fairly straightforward to implement.</p>								
DECnet implementation	<p>From the start we wanted to use the EARN/ED provided so-called “GBOXes” for DECnet routing. It took several months to get the administrative work settled. The GBOXes were used for evaluation of the EARN NJE/OSI and other solutions to carry the NJE traffic. Before the true operational phase, which started as late as February 1990, DECnet routing was done using the cisco routers. Operational experience has not been great. Using DECnet over Ethernet is not optimal. The users experienced “time-outs” and broken file transfers due to either congested links or bridges. Both DEC and Vitalink have been helping out in analyzing these problems.</p>								

continued on next page

Profile: NORDUnet (*continued*)

EARN/NJE

The first solution to carry the traffic over the new links was by using band-splitting. The evaluation of EARN NJE/OSI and other technical possibilities was ready late 1989. The results are documented in a report by Harri Salminen. Currently a mix of solutions are in place. NJE/DECnet (JNET), NJE/IP (VMNET), and NJE/OSI/X.25 are all in use. The future solution depends to some extent on the strategy decided on by EARN, but a more homogeneous solution would be preferable.

X.25 experience

X.25 is the most extensive undertaking in the whole implementation project. The evaluation of possible equipment pointed us to a switch manufactured by Sattelcom in France. Originally it only supported LLC1. The error recovery on packet level was not fully implemented either. With our situation of noisy lines and possible packet losses at the bridges, this resulted in hanging connections.

The development of LLC2 (X.25 over LAN) improved this situation. With a joint effort between HEPnet, EUnet/CWI and RUNIT in Norway, management and statistics software has been developed. X.25 has now started becoming a full service. The COSINE IXI X.25 pilot service will make possible X.25 connectivity to a set of new countries which are hard to reach with existing services.

Current activities

Before the summer of 1990, Iceland was connected by a leased 9.6Kbps line to KTH/Stockholm. The US connection was planned changed from the current satellite 56Kbps to a terrestrial 64 or 128Kbps circuit, still maintained by JvNC. A longer term solution is currently being discussed with a set of European partners to be able to reach a common "fat-pipe" solution.

The main area of work is under the RIPE initiative. Other activities include collaboration with the IBM/EASI initiative to strengthen the bandwidth to central Europe and possibly to the US.

The COSINE RARE IXI pilot X.25 network will be used for OSI pilots and other services. Most likely IP connections to Spain and Ireland will be set up.

The main emphasis for 1990 is on the following activities:

ISO IP DECnet Phase V pilot: This project aims at providing a pilot service for ISO IP and DECnet Phase V. The intention is to test inter-working between different equipment providing ES-ES and ES-IS routing. Other items for study is NSPS allocation, routing and administrative domains. Of interest is also DECnet Phase V migration. The experience gained from this pilot will help us in planning for a full-blown ISO CLNS service as well as introduction of DECnet Phase V in NORDUnet and the national networks.

Introduction of "NETF," The NORDUNET Engineering Task Force: The NORDUNET Board has decided to form a new forum for the technical work. The model is the Internet Engineering Task Force, IETF. The first NETF meeting will be in the autumn 1990.

X.500 experiments: NORDUNET is active in the X.500 area. Some five to six DSAs have been active in the R&D pilot/experiment informally coordinated by University College London.

Future plans

The current NORDUNET program will end after 1991. Finding longer term financing for development and engineering in the networking field will be a major activity during 1990-91. Another major field of work is the introduction of OSI services in general. NORDUNET will most likely work in that area for several years to come.

New technologies such as high speed networking will be of increased importance during the 1990s. A follow-up technical solution to the current Ethernet MAC level system used in NORDUNet will be investigated. The PTT's plans for higher speed services will play an important role in this planning.

Contact information**NORDUNET**

c/o SICS PO Box 1263

S-164 28 Kista

Sweden

Phone: + 46 8 752 1563

Fax: + 46 8 751 7230

E-mail: NORDUNET@sics.se

MATS BRUNELL studied Computer Science at the University of Stockholm. He was responsible for the Department of Computer Planning at the University between 1981 and 1983. From 1983 to 1986 he was responsible for the Communications group at the University Computer Center. Since 1986 he has been working on the NORDUNET project team. Throughout the 80s he has been actively involved in International research networking activities, such as SUNET, the Swedish University Network, EARN and RARE activities.

Clark and Kleinrock receive 1990 SIGCOMM Award

The ACM Special Interest Group on Data Communication (SIGCOMM) announced on September 25, 1990 that it had given its annual SIGCOMM Award to Dr. Leonard Kleinrock and Dr. David D. Clark. Dr. Kleinrock was recognized for "his seminal role in developing methods for analyzing packet network technology." Dr. Clark was recognized for "major contributions to internet protocols and architecture." The Awards were presented at the annual SIGCOMM conference.

The SIGCOMM Award is a "lifetime" achievement award, given annually to a person who has a history of important contributions to the field of computer communication. This year the award committee had a tie vote and decided to recognize both of the top two candidates.

Dr. Kleinrock is a professor at UCLA. He is an IEEE Fellow and a member of the National Academy of Engineering. He is best known for his classic textbooks on queuing theory and on computer networks.

Dr. Clark is a senior research scientist at the MIT Laboratory for Computer Science. He was chairman of the Internet Activities Board, the policy and standards board for the TCP/IP protocol suite, through much of the period during which TCP/IP was developed.

ACM SIGCOMM is a professional society for persons in the field of computer communications and was established in 1969, as a special interest group of the Association for Computing Machinery. It currently has about 5,000 members. The annual ACM SIGCOMM conference is one of the major research gatherings in computer communication. This year's conference was held at the University of Pennsylvania, and drew approximately 300 people. The conference received support from the FTP Software Inc., Digital Equipment Corporation's Network Systems Lab, Interop Inc., Bellcore, and Advanced Computer Communications.

Press Here for The Internet: Let's Promote More Good News!

by Daniel P. Dern

Background

For INTEROP 90, I ran a BOF (Birds-Of-a-Feather session) on the topic of "Getting (Better) Press for the Internet, TCP/IP and UNIX," encompassing some combination of the net per se, the community surrounding it, the the associated technologies, vendors, users, etc. We got an interesting mix of press and users. (Next year, I'd also like to see a few PR people in attendance.)

The BOF, combined with other interactions I had during the week of the show, confirmed my sense of how things do (and don't) work, to wit: there is still a large gap in understanding between the Internet community and the rest of the world. All too often, neither side clearly understands what the other is all about—and it is incumbent on the Internet community to close the gap.

Examples: "Why doesn't the press call the academic networks and computing sites?" "I thought INTEROP began as a LAN show." "Is there anything here relevant to Macintosh users?"

Here's some of what we talked about:

During the past several years, the Internet has "made the news" in the trade, business and general press, and even onto radio and television, undoubtedly more than in the previous two decades.

Top stories

Unfortunately, the topics have been predominantly less than felicitous. The big Internet news stories? The *Internet Worm*, of course, reprised by coverage and editorializing the trial. Cliff Stoll's stalking of the *Wiley Hacker* probably probably takes second place. With the "newsworthiness" ice broken, the past year or so then gave us the NASA WANK Worm, several cases of Mailing List Madness, and the factoring of the superprime number.

There has been some pro-net news coverage, like the *Gore Supercomputing Bill* and the NSFNET backbone—and even that last one got its tail twisted a tad. The work in the *Internet Activities Board* (IAB) and *Internet Engineering Task Force* (IETF) gets regular, if short, coverage. Internet "names" like Vint Cerf, Marshall Rose, Phill Gross, Stephen Wolff, and Jim Herman get quoted frequently. And Internet-relevant stories like TCP/IP, UNIX, OSI/GOSIP, routers, and vendor/product announcements help keep the Internet in the corner of the public eye.

But much of the press and its readers remain unaware of what's happening in Internet-land—and changing that is an important part of the growth of the Internet, and its many offspring of networks, companies and technologies.

Benefits

Getting more Internet-related press coverage is likely to be an all-around "win":

- More press mentions for your organization, in an increasingly competitive scene.
- More positive coverage of Internet issues.
- You get in the news—good for your career, one hopes.

What you can do

What can members of the Internet community do to promote the various causes? Here's a few suggestions:

- Work *with*, not *around*, the PR department/agency: If somebody is responsible for working with the press, work through them. Get their (and your manager's) approval. Let the PR guys'n'gals do their job, it will save you time. (And you'll last longer if you aren't a loose cannon.) If it works out, they'll come back for more.
- Identify "good news" to promote: Are you involved in some Internet-related activity? Are your new products important to Internet users? Are you using some Internet vendor's product or service in your work? Help point out "success stories" in terms of specific applications, identify topics that may result in press coverage. (Be prepared to spend some time educating PR people about the Internet and why it's important.)
- Be a resource: offer to write articles and to be available to the press. [Ed.: I like this one, especially if you include *ConneXions* in the definition of "press"]

Guidelines

When working with the press, follow standard guidelines, e.g.:

- Never say anything you aren't ready to see in print.
- Get clearance from PR and your manager, as needed.
- Return phone calls promptly—reporters are usually on a tight deadline.
- Keep explanations focused and simple, but be willing to provide explanations.
- Graphics, maps, charts or photographs may improve your chances for coverage.

Be nice to them!

And spend some time learning about the press—after all, you're asking them to make a similar effort. Read the publications you think are most relevant.

Identify four or five editors and reporters who cover areas relevant to you, and where what you do is relevant. Call or e-mail; offer to be a resource in terms of offering background expertise, quotes, perhaps a source of potential stories.

If you're at conferences, trade shows and other events, introduce yourself to the press to establish a working relationship with editors and reporters. Remember—press people need willing information sources and newsworthy subjects!

Next year

I'd like to do this session again next year, and see how much progress we've made. But even more important, I'm hoping that over the coming year, more people will help close the Internet "communication gap," and help generate more, positive, relevant press coverage for all the world-changing excitement that continues to flow from the Internet.

DANIEL P. DERN is a Watertown, Mass-based free-lance writer specializing in technology, science and industry. A frequent contributor to *ConneXions*, including last year's ARPANET historical retrospective, Dern writes for leading publications and vendors in the network and computer industry, as well as writing hi-tech humor columns, science fiction, musical theater, and teaching UNIX. He was previously PR Manager at BBN Communications. He can be reached via e-mail as ddern@world.std.com (Internet) or dandern (MCIMail).

Book Reviews

The *User's Directory of Computer Networks*, edited by Tracy Lynne LaQuey. Published by Digital Press, ISBN 1-55558-047-5, 1990. Also registered with a Prentice-Hall ISBN: 0-13-950262-9.

Introduction

Today's widespread analogy that likens computer networking to the highway system logically leads to the observation, made by Tracy LaQuey, that the network traveler needs a *roadmap* to get around. She intends the User's Directory of Computer Networks to be the tool that helps network users understand the communications paths, see how they connect, locate resources (machines, services, or people) that they need, and understand some basic networking concepts.

Organization

The Directory is a descendant of a 1987 volume of the same title published by the University of Texas, and edited by Carol Englehardt Kroll, which was subsequently revised by LaQuey. The current directory is divided into chapters that discuss specific networks, such as the DECnet Internet and the Internet, essays on the *Domain Name System*, the *OSI Directory Service* (X.500), Electronic Mail, and an organizational index. This volume, includes several more networks and has an expanded narrative about each network.

Content

This book successfully pulls together a lot of information in a consistent and coherent presentation. Most chapters (several of which have subchapters that describe component networks of an internet) provide descriptions that answer the same key questions about each network: what is the topology? what protocols are supported? what services are provided? what are the membership requirements? how is the network administered? what are the usage guidelines? The descriptions don't go into great technical depth, but that's not the editor's goal. The directory provides maps and extensive lists of hosts, contacts, and network numbers for reference purposes, but the reader comes away from a chapter chiefly with a useful overview of each network and a basic understanding of where each fits into the big picture.

The essays in the final chapters are particularly helpful for users who have a limited amount of networking experience. John Quarterman presents a good summary of the complex issues of electronic mail, and provides a bibliography for those readers who want a more extensive treatment of email. Mic Kaczmarczik includes a useful set of tables designed to help users construct and send messages between many of the networks described in the directory.

Paul Mockapetris contributed the chapter on Domains. In a succinct three-and-a-half pages, he does a neat job of summarizing the important concepts of the Domain Name System and describing why the reader should care about them. A list of domain names is included.

The OSI X.500 chapter [reprinted from the June 1989 issue of *ConneXions*] contains more detail than the other essays and is less conversational in tone. The focus here is more on the technical specifics of the OSI Directory and is aimed at a more technically sophisticated audience.

The final chapter, List of Organizations, is a valuable cross-reference that gives the reader a picture of the connectivity of over five thousand organizations.

A shortcoming of the Directory is one that is typical of all books dealing with an area that is developing as quickly as networking—some percentage of the data is automatically outdated as soon as the text is given to the publisher.

However, even if specifics change over time, such as contact names, the information that remains serves as a starting point for finding the most current information.

Recommended

The User's Directory is impressive for several reasons. It presents a huge quantity of information in a straightforward and comprehensible way. Tracy has done an excellent job of editing that doesn't make the user feel overwhelmed by a subject that can actually be quite overwhelming to those not immersed in network technology. Her efforts at collecting and verifying information are apparent, and her diligence proves worthwhile. This reference guide will occupy a prominent place on the bookshelves of the masses of network users who need the information that LaQuey has compiled.

—Karen Roubicek

The Simple Book: An Introduction to Management of TCP/IP-based internets, Marshall T. Rose, ISBN 0-13-812611-9, Prentice-Hall, 1990.

Marshall Rose's latest book, *The Simple Book: An Introduction to Management of TCP/IP-based internets*, continues his elucidation of networking technology. With the rapid growth of computer networks in education, government and industry, the topic of *network management* becomes more compelling. There are now more networking devices in the marketplace than there have ever been before, and furthermore, the future does not appear to offer a respite. With this explosion of internetworking products, their management becomes all the more importunate.

Network management seems to evoke different opinions from everyone. This is probably due to its recent emergence as a technology. Marshall's book offers to light up the subject based on real-world operational experience and some very sharp insight. Like his last work, *The Open Book*, he includes a number of "soap boxes" where he offers his unique opinion on matters both technical and political. There is much room for opinions in this new field of network management, but Marshall endeavors to guide you down a well-lit path.

History

The book opens with a brief history of network management in the Internet. Missing is any coverage of the proprietary techniques used to monitor or manage internetworks before the history section began, but the history that Marshall does report is quite colorful. He covers how he got involved in network management and opens the door for a brief look at some of the major issues.

I was a bit surprised to find an introduction to OSI; however, it is one of the better explanations of the basic structure and design of OSI and serves to lay the groundwork for later discussions. One of the more useful topics is the OSI framework for network management. Rose explains some of the thinking behind the design of CMIS/CMIP.

The TCP/IP Protocol Suite

The OSI framework is followed by an introduction to the TCP/IP protocol suite. Included in this chapter is an explanation of how the TCP/IP standardization process works. As far as I am aware, this is the first time this process has appeared in print outside of RFCs or verbal discussions. The TCP/IP architectural model is discussed with a section devoted to each of the 4 layers; *interface*, *internet*, *transport* and *application*. These sections comprise a very good and complete précis of the TCP/IP protocol suite. From ARP to TCP the language is clear, concise and accurate. The section on routing covers subnetting, EGP and even mentions BGP.

continued on next page

Book Reviews (*continued*)

The last part of the chapter covers *Ping*, *Traceroute* and network monitoring. This chapter stands out as one of the better introductions to the TCP/IP architecture. It alone makes the book worth reading.

Fundamental Axiom

The third chapter starts the real network management discussions and offers the first hint behind the book's title. Given the number and diverse type of internetworking products, Marshall defines his fundamental axiom:

The impact of adding network management to managed nodes must be minimal, reflecting a lowest common denominator.

Thus begins the explanation of why there is the *Simple* in SNMP. ASN.1 basics for SNMP is introduced. Anyone who has implemented ASN.1 from the ISO specifications will appreciate this section. Details of the *Basic Encoding Rules* (BER) are covered later.

SMI and MIB

The *Internet-standard Network Management Framework* is introduced with the *Structure of Management Information* (SMI) and *Management Information Base* (MIB) documents. The SMI is covered first. Much of what is discussed is not presented together this cohesively anywhere else. Both MIB I and MIB II are explained with the differences highlighted. Case diagrams are introduced and used to help demonstrate how MIB variables relate to each other. This section will be invaluable to a potential implementor as a few points are covered that are not mentioned anywhere in the RFCs!

SNMP

Chapter 5 discusses the *Simple Network Management Protocol* (SNMP). Here is where Marshall ties together everything covered so far. It also contains the best definition of "community" I have seen. The one nit I have to pick is with the section on "Row Addition." Rose explains that a row can be added to a table by performing a SET operation on all of the columns in that row with the same instance identifier. What happens if one of the columns in a row is omitted?

Interestingly enough, Marshall admits that the hardest part of the network management problem is writing the *applications* that use the SMIs, MIBs and protocols. A brief look at a set of applications is presented along with some examples on how to use them to solve networking problems.

Implementations

Marshall has written a complete SNMP agent for 4BSD UNIX which he provides complete with his *ISO Development Environment* (ISODE) package. Ordering information for ISODE can be found in an appendix. Included is a tool for rapid-prototyping network management applications. Both are covered in some detail including data structures and support subroutines to help clarify the design. Also included is a description of a mechanism for distributing MIB support among multiple UNIX processes. This new protocol is called *SNMP Multiplexing* (SMUX).

The Final Soap Box

Marshall offers his vision of the future in the last chapter as one big soap box. I was a bit surprised by the lack of soap boxes throughout the text but found he saved them for the final chapter.

Read it!

This one book contains information that will help both beginner and expert. If you do not know what network management is and want to find out, then read it. If you do know what network management is but have some questions or want to write a MIB, then read it. Simply put, read it!

—Greg Satz

Upcoming Events

The *Second International Symposium on Integrated Network Management* will be held April 1-5, 1991 at the Marriott Crystal Gateway, Crystal City across the river from Washington, DC.

This year the International Federation for Information Processing (IFIP) Working Group 6.6 is teaming with the Institute of Electrical and Electronics Engineers (IEEE) Communication Society Technical Committee on Network Operations & Management (CNOM) to broaden the scope of the Symposium.

Like the first Symposium, held in Boston in 1989, this meeting will create a forum for information exchange and cooperation among vendors, system integrators, researchers, standards developers, and users. The themes of the Second International Symposium on Integrated Network Management grow out of the 1989 Symposium, which highlighted the increasing need for Total Systems Management to integrate all aspects of managing data communications, telecommunications, distributed systems and distributed enterprise applications.

Topics

Authors have been invited to submit unpublished papers, as well as proposals for tutorials, panel discussions or informal (birds-of-feather) sessions in the following topic areas:

- User Requirements for Total Systems Management
- Models and Architectures
- Standards Issues: OSI, TMN, Internet and Other Standards
- Fault, Configuration, Accounting, Performance & Security Mgmt.
- Quality of Service Management
- Telecommunications Management, Including OAM&P
- Management Information, Definition and Storage
- Management Protocols
- Management of Heterogeneous Systems
- Management Domains—Principles and Practice
- Securing Management Systems
- Interplay of Security and Management
- Inter-Organizational Management Issues
- Artificial Intelligence Techniques
- User Interfaces and Management Languages
- Application of Distributed Computing Technologies to Management
- Interplay of Distributed Operating Systems and Management
- Implementations, Case Studies, Lessons Learned, Horror Stories
- Other Related Topics

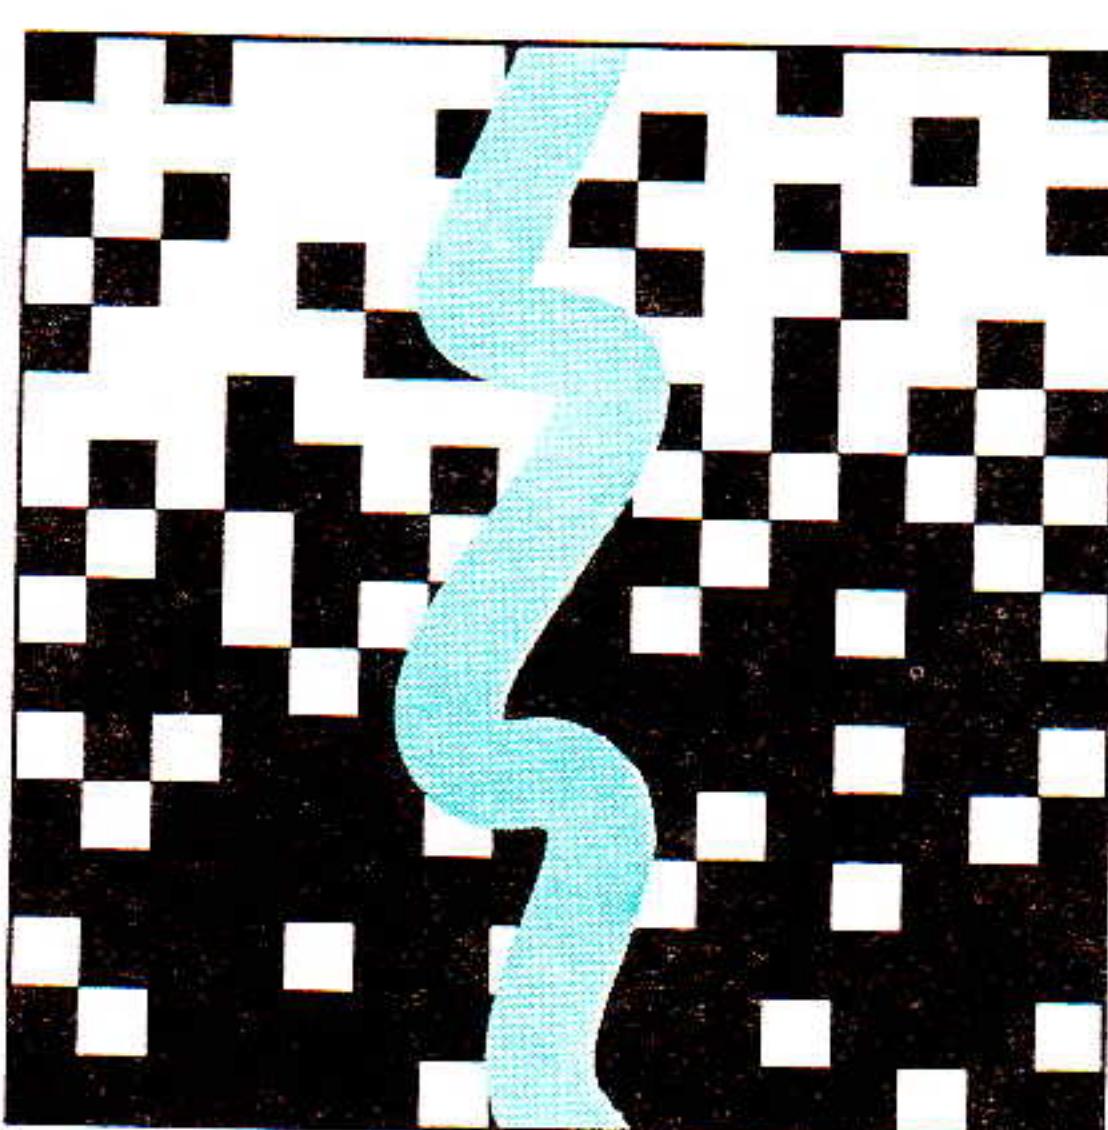
Vendor participation

Vendors interested in becoming Symposium Patrons or Friends and being given the opportunity to present products and/or plans should contact:

Ms. Kimberly Kappel,
Georgia Institute of Technology
404-853-9383
kappel@prism.gatech.edu

For more information about the symposium contact:

*Second International Symposium
on Integrated Network Management*
P.O. Box 191885
San Francisco, CA 94119-1885
415-392-3751.



GOSIP Document order information

by Richard Colella, NIST

This article contains information needed to obtain the U.S. *GOSIP* and *NIST/OSI Implementors Workshop* (OIW) documents.

GOSIP

GOSIP Version 1: This document (FIPS 146) was published in August 1988. It became mandatory in applicable federal procurements in August 1990. The NIST Point of Contact is Jerry Mulvenna. The order number is: FIPS PUB 146 The document is also available on-line from host `nic.ddn.mil` as: <PROTOCOLS>GOSIP-V1.TXT

GOSIP Version 2: This document was issued as a draft early in 1990. It has undergone public review and comment. Version 2 was revised to address the comments received from industry and government; the revisions were subsequently reviewed and approved by the *GOSIP Advanced Requirements Committee* in March, 1990. Version 2 is expected to become a Federal Information Processing Standard (FIPS) by the end of October, 1990. The NIST Point of Contact for this document is Jerry Mulvenna. Hardcopy can be obtained from the NIST Standards Processing Coordinator (ADP) or on line from Internet host `nic.ddn.mil`: <PROTOCOLS>GOSIP-V2-DRAFT.DOC

Implementation Agreements

The output of the NIST Workshop for Implementors of OSI (OIW) is a pair of aligned documents, one representing *Stable Implementation Agreements* (SIA), the other containing *Working Implementation Agreements* (WIA) that have not yet gone into the stable document. Material is in either one or the other of these documents, but not both, and the documents have the same index structure.

The SIA is reproduced in its entirety at the beginning of each calendar year, with an incremented version number. Replacement page sets are distributed subsequently three times during each year (after each Workshop), reflecting errata to the stable material. The replacement pages constitute the next edition of that year's version.

The WIA is reproduced in its entirety after each Workshop (held in March, June, September and December). OIW attendees automatically receive the WIA. The NIST Points of Contact for the OIW are Tim Boland, and Brenda Gray.

SIA, Version 1, Edition 1 (Dec., 1987): is published as NIST Special Publication 500-150. It is the appropriate version and edition of the SIA for GOSIP Version 1 (FIPS 146). Hardcopy can be obtained from the U.S. Government Printing Office. The GPO Stock Number is 003-02838-0. It is also available from NTIS, Order Number: PB 88-168331.

SIA, Version 1, Edition 3 (August, 1988): is also published as NBS Special Publication 500-150 and available from the U.S. Government Printing Office. (note the different GPO Stock Number when ordering). GPO Stock Number: 003-003-02838-0. You can also get the SIA via FTP from `nic.ddn.mil`: <PROTOCOLS>NBSOSI-AGREEMENTS.DOC

SIA, Version 2, Edition 1 (Dec., 1988): is published as NBS Special Publication 500-162. The GPO Stock Number is 003-003-02921-1. Also available from the IEEE Computer Society Press, ISBN 0-8186-9022-4, Catalog No. 2022, or from NTIS, Order Number: PB 89193312.

SIA, Version 2, Editions 2-4: These are available as hardcopy from NIST staff, subject to staff availability. Contact Brenda Gray, NIST.

SIA, Version 3: Beginning with SIA, Version 3 (NIST Special Publication 500-177), the GPO is offering a subscription service for the Stable Agreements. This comes in two flavors. First, you may purchase the base document (V3E1) and the three sets of change pages will be sent to you automatically when they become available to GPO (Note: each set of change pages was previously considered a new "edition." This terminology is no longer in use.). Second, those who have purchased V3E1 separately may pick up the subscription service beginning with the June change pages at a reduced charge. This option is intended for use only during this first year of the subscription service. Note that the GPO stock number is the same for both options. GPO Stock Number: 903-015-00000-4.

SIA, Version 3, Edition 1 (Dec., 1989): is published as NBS Special Publication 500-177. NTIS, Order Number PB 90-212192 or IEEE Computer Society Press, ISBN 0-8186-2075-7, Catalog No. 2075.

WIA (March, 1990): published as a NIST Interagency Report (NISTIR-4302), is the most recent copy of the WIA. Available from NTIS, Order Number NISTIR-4302.

GOSIP User's Guide

GOSIP Users' Guide: This publication assists federal agencies in planning for and procuring OSI. It provides tutorial information on OSI protocols as well as information on OSI registration, GOSIP technical evaluation, and GOSIP transition strategies. Available from NTIS, Order Number PB 90-111212.

Contact information

Jerry Mulvenna mulvenna@ecf.ncsl.nist.gov
Technology, B217
Gaithersburg, MD 20899
301-975-3631

Tim Boland boland@ecf.ncsl.nist.gov
Chairman, OIW
Technology, B217
Gaithersburg, MD 20899
301-975-3608

Brenda Gray, OIW Registrar
Technology, B217
Gaithersburg, MD 20899
301-975-3664

National Technical Information Service (NTIS)
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161
703-487-4650

IEEE Computer Society
Press Order Department
10662 Los Vaqueros Circle
Los Alamitos, CA 90720
1-800-272-6657

U.S. Government Printing Office
Washington, DC 20402
202-783-3238

Standards Processing Coordinator (ADP)
National Institute of Standards and Technology
Technology Building, Room B-64
Gaithersburg, MD 20899
301-975-2816

CONNE~~X~~IONS
 480 San Antonio Road
 Suite 100
 Mountain View, CA 94040
 415-941-3399
 FAX: 415-949-1779

Bulk Rate
 U.S. POSTAGE
 P A I D
 SAN JOSE, CA
 PERMIT NO. 1

ADDRESS CORRECTION
 REQUESTED

CONNE~~X~~IONS

EDITOR and PUBLISHER **Ole J. Jacobsen**

EDITORIAL ADVISORY BOARD **Dr. Vinton G. Cerf, Vice President,
 Corporation for National Research Initiatives.**

**A. Lyman Chapin, Senior Consulting Engineer,
 Data General Corporation.**

**Dr. David D. Clark, Senior Research Scientist,
 Massachusetts Institute of Technology.**

**Dr. David L. Mills, Professor,
 University of Delaware.**

**Dr. Jonathan B. Postel, Communications Division Director,
 University of Southern California,
 Information Sciences Institute.**

Subscribe to CONNE~~X~~IONS

U.S./Canada \$125. for 12 issues/year \$225. for 24 issues/two years \$300. for 36 issues/three years

International \$ 50. additional per year **(Please apply to all of the above.)**

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

Check enclosed (in U.S. dollars made payable to **CONNE~~X~~IONS**).

Visa MasterCard American Express Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to: CONNE~~X~~IONS

Back issues available upon request \$15./each
 Volume discounts available upon request

480 San Antonio Road Suite 100
 Mountain View, CA 94040
 415-941-3399 FAX: 415-949-1779

CONNE~~X~~IONS